




Solenoidal Detector Collaboration



TECHNICAL DESIGN REPORT

1 April 1992

SOLENOIDAL DETECTOR COLLABORATION

TECHNICAL DESIGN OF A DETECTOR

to be operated

at the

Superconducting Super Collider

1 April 1992

Contact person:

George H. Trilling, Spokesperson
50A-2160 Lawrence Berkeley Laboratory
Berkeley CA 94720
510/486-6801, Bitnet GHT@LBL

Members of the Solenoidal Detector Collaboration

Argonne National Laboratory: E. L. Berger, R. E. Blair, J. W. Dawson, M. Derrick, T. H. Fields, V. Guarino, N. F. Hill, P. K. Job, T. B. W. Kirk, E. N. May, J. Nasiatka, L. J. Nodulman, L. E. Price, J. Proudfoot, H. M. Spinka, R. L. Talaga, D. G. Underwood, R. G. Wagner, A. B. Wicklund

University of Arizona: K. A. Johns

Beijing University, China: He Yu Ming, Lai Chu Xi, Liu Hong Tao, Liu Song Qiu, Lou Bing Qiao, Yang Ji Xiang, Yao Shu De, Zhang Re Ju

Institute of High Energy Physics, Beijing: H. Cui, C. Gao, M. Gao, W. Gao, S. Gu, H. Hu, D. Huang, W. Li, N. Liu H. Mao, H. Ni, N. Qi, G. Rong, H. Shi, S. Tang, P. Wang, T. Wang, X. Xie, S. Xue, W. Yan, M. Ye, Ge. Zhang, W. Zhao, L. Zheng, Y. Zhou, Q. Zhu, Y. Zhu

Brandeis University: S. Behrends, J. R. Bensinger, P. Kesten, L. Kirsch

Bratislava State University, Czechoslovakia: P. Povinec, P. Strmen

University of Bristol: B. Foster, G. P. Heath

Brown University: D. Cutts, R. Hlustick, R. Partridge

Institute for Physics and Nuclear Engineering, Bucharest, Romania: A. Alexa, M. Horoi, D. Pantea, M. Pentia, C. Petrescu

California Institute of Technology: A. J. Weinstein

University of California at Davis: J. Gunion, D. Pellett, S. Mani

University of California at Irvine: A. Lankford, B. Schmid, D. Stoker

University of California at Los Angeles: K. Arisaka, H.-U. Bengtsson, C. Buchanan, D. Chrisman, D. Cline, J. Hauser, T. Muller, J. Park, D. Roberts, W. Slater

University of California at Riverside: J. Ellison, A. Heinson, A. Klatchko, C. Lietzke, S. J. Wimpenny

University of California at Santa Cruz: J. DeWitt, D. Dorfman, C. Heusch, B. Hubbard, D. Hutchinson, A. L. Litke, W. Lockman, K. O'Shaughnessy, W. Rowe, H. Sadrozinski, A. Seiden, E. Spencer, A. Webster

Carleton University: J. Armitage, P. Estabrooks, S. Godfrey, M. O'Neill

Chiba University: H. Kawai

University of Chicago: C. Campagnari, M. Contreras, S. Eno, H. Frisch, C. Grosso-Pilcher, M. Miller, L. Rosenberg, H. Sanders, M. Shochet, G. Sullivan

University of Colorado: G. J. Baranko, B. Broomer, H. W. K. Cheung, J. P. Cumalat, M. Daoudi, E. Erdos, W. T. Ford, S. V. Greene, D. Johnson, U. Nauenberg, P. Rankin, G. Schultz, V. Slonim, J. G. Smith

Colorado State University: W. Toki

Joint Institute for Nuclear Research, Dubna: V. I. Astakhov, B. V. Batyunia, A. Bischoff, Y. A. Budagov, A. M. Chuenko, A. I. Dokshin, M. Finger, S. B. Gerasimov, V. M. Golovatyuk, Yu. N. Gotra, Yu. S. Gusar, Z. Guzik, D. I. Hubua, Yu. V. Ilyin, S. V. Kashigin, Y. N. Kharzheev, I. F. Kolpakov, A. D. Kovalenko, F. V. Levchanovsky, Y. F. Lomakin, A. I. Malakhov, E. A. Matyushevsky, A. A. Omelianenko, Yu. S. Pakhmutoy, Y. A. Panebratsev, I. V. Puzynin, A. A. Semenov, A. E. Senner, A. V. Shabunov, V. T. Sidorov, A. N. Sinaev, A. N. Sissakian, V. A. Smirnov, T. Spassoff, E. N. Tsyganov, I. A. Tyapkin, L. A. Vasilyev, G. V. Velez, V. B. Vinogradov, A. S. Vodopianov, V. Vrba, Y. V. Zanevsky, N. I. Zhuravlev, N. I. Zimin, A. I. Zinchenko

Duke University: A. T. Goshaw, S. H. Oh, T. J. Phillips, W. J. Robertson, C. Wang

Erevan Institute of Physics: A. C. Amatuni, G. A. Vartgapetian

Fermilab: V. H. Areti, M. Atac, E. Barsotti, L. Bartoszek, A. E. Baumbaugh, A. Beretvas, R. Bernstein, M. Binkley, A. D. Bross, A. Byon-Wagner, C. Campagneri, S. Cihangir, J. W. Cooper, B. Denby, K. G. Denisenko, N. L. Denisenko, T. Droege, D. P. Eartly, J. E. Elias, R. W. Fast, D. Finley,

G. W. Foster, J. Freeman, I. Gaines, S. A. Gourlay, D. R. Green, J. Grimson, C. Grozis, S. R. Hahn, R. M. Harris, J. Hoff, J. Huth, J. Hylen, U. Joshi, R. D. Kephart, J. Kilmer, D. H. Kim, H. J. Krebs, A. Lee, P. J. Limon, R. Luther, P. Mantsch, F. Markley, P. S. Martin, A. Mukherjee, T. Nash, C. Newman-Holmes, A. Para, J. Patrick, R. Plunkett, E. E. Schmidt, S. L. Segler, R. P. Stanek, A. Stefanik, J. Strait, H. J. Stredde, S. Tkaczyk, R. Vidal, R. L. Wagner, R. H. Wands, R. Yarema, G. P. Yeh, J. Yoh, T. Zimmerman

University of Florida: R. Field, J. Harmon, J. Walker

Florida State University: M. Corden, J. H. Goldman, V. Hagopian, K. Johnson, H. Wahl, S. Wasserbach, D. Xiao, S. Youssef

Fukui University: M. Kawaguti

Gomel State University: A. M. Dvornik, N. B. Maksimenko

Harvard University: G. Brandenburg, J. Carnevale, G. Feldman, M. Franklin, J. Konigsberg, J. Oliver, E. Sadowski, P. Schlabach, R. Wilson

University of Hawaii: C. Kenney, S. Parker

Hiroshima University: Y. Iwata, T. Ohsugi, A. Sakaguchi

Hiroshima Institute of Technology: M. Asai

Ibaraki College of Technology: M. Shiiden

University of Illinois at Chicago: M. Chung, H. Goldberg, S. Margulies, H. Mendez, J. Solomon, F. Vaca

University of Illinois at Urbana: R. Downing, S. Errede, A. Gauthier, M. Haney, L. Holloway, I. Karliner, A. Liss, T. O'Halloran, J. Thaler, P. Sheldon, V. Simaitis, J. Wiss

Indiana University: B. Brabson, T. Collins, A. Dzierba, G. Hanson, X. Lou, F. Luehring, T. Marshall, B. Martin, H. Ogren, D. Rust, E. Wente, A. Zieminski, D. Zieminski

Iowa State University: W. Anderson, P. Dervenagas, J. Dyke, J. Hauptman, M. Pang, M. Wendling, A. Zinchenko

Johns Hopkins University: J. A. Bagger, B. A. Barnett, B. J. Blumenfeld, J. F. Krizmanic, J. Skarha, F. Snider

National Laboratory for High Energy Physics (KEK), Japan: F. Abe, K. Amako, Y. Arai, Y. Doi, H. Fujii, Y. Fukui, T. Haruyama, H. Ikeda, S. Inaba, T. Inagaki, H. Iwasaki, S. Kabe, N. Kanematsu, J. Kanzaki, T. Kondo, A. Maki, A. Manabe, M. Mishina, Y. Morita, M. Nomachi, S. Odaka, K. Ogawa, T. K. Ohoka, H. Sakamoto, O. Sasaki, Y. Takaiwa, S. Terada, T. Tsuboyama, K. Tsukada, N. Ujiie, Y. Unno, Y. Watase, A. Yamamoto, Y. Yasu

Slovak Academy of Science, Košice, Czechoslovakia: F. Krivan, M. Seman, J. Špalek

University of Kansas: R. Ammar, S. Ball, P. Baringer, D. Coppage, R. Davis, N. Kwak

Kyoto University: R. Kikuchi, K. Miyake

Lawrence Berkeley Laboratory: G. S. Abrams, H. Aihara, A. Barbaro-Galtieri, R. M. Barnett, L. P. Bautz, C. J. Bebek, R. N. Cahn, A. Ciocio, C. A. Corradi, C. Day, S. Dow, A. Dubois, P. H. Eberhard, W. R. Edwards, K. Einsweiler, T. Elioff, R. Ely, G. Gabor, M. G. D. Gilchriese, W. Greiman, D. E. Groom, C. Haber, C. Hearty, I. Hinchliffe, M. Hoff, R. Jared, R. W. Kadel, J. A. Kadyk, F. Kirsten, I. Kipnis, S. Kleinfelder, M. E. Levi, A. Lim S. C. Loken, N. Madden, O. Milgrome, J. Millaud, Y. Y. Minamihara, T. L. Moore, D. R. Nygren, E. Oberst, W. L. Pope, M. Pripstein, J. Rasson, D. Santos, B. Schumm, M. Shapiro, D. Shuman, H. G. Spieler, R. Stone, M. Strovink, W. Thur, G. H. Trilling, R. C. Weidenbach, W. A. Wenzel, S. Wonduke, M. Wong

University of Liverpool: T.J.V. Bowcock, J. B. Dainton, E. Gabathuler, T.J. Jones, S. J. Maxfield

Los Alamos National Laboratory: G. Dransfield, M. Gamble, M. Gold, C. Grastataro, J. Hanlon, J. A. J. Matthews, F. Michaud, W. Miller, A. P. T. Palounek, R. Reid, R. Ricketts, C. Sadler, W. Sailor, S. Seidel, R. Shea, D. Skinner, T. Thompson, J. Williams K. Woloshun, H. Ziock

Louisiana State University: A. Fazely, R. Imlay, S. Khosravi, C. Lyndon, G. McMills, W. Metcalf
University of Manitoba: G. Smith
University of Maryland: A. R. Baden, A. H. Ball, C. Y. Chang, D. G. Fong, J. A. Goodman, N. J. Hadley, A. Jawahery, R. G. Kellogg, S. Kunori, A. Skuja, G. T. Zorn
McGill University: K. Ragan, D. G. Stairs
University of Michigan: D. Amidei, R. C. Ball, M. Campbell, J. Chapman, K. De, P. Derwent, H. R. Gustafson, K. Hashim, S. Hong, L. W. Jones, S. B. Kim, M. J. Longo, J. Mann, M. R. Marcin, H. A. Neal, D. Nitz, B. P. Roe, G. Snow, R. Thun, D. Wu, S. Zhang
I.N.F.N., Milano: G. Cesura, P. Inzani, P. F. Manfredi, D. Menasce, L. Moroni, D. Pedrini, V. Re, S. Sala, F. Svelto
University of Minnesota: P. Border, H. Courant, R. Gray, K. Heller, Y. Kubota, M. Marshak, E. Peterson, R. Poling, K. Ruddick
Academy of Science of BSSR, Minsk: J. A. Kulchitsky, L. G. Moroz
University of Mississippi: B. Moore, D. Summers
Miyazaki University: T. Nakamura
University of Montreal: L.-A. Hamel, J.-P. Martin, L. Martinu, P. Taras, M. Wertheimer
Institute of Theoretical and Experimental Physics, Moscow: B. Ejov, Ju. Grishkin, V. Homyakov, V. Kozitsyn, V. Krylov, O. Pogorelko, P. Shyshov, O. Zemskov
Nagoya University: M. Nakamura, K. Niwa
Naruto University of Education: H. Yoshida
Niigata University: K. Miyano, H. Miyata
Northeastern University: G. O. Alverson, M. J. Glaubman, I. D. Leedom, J. H. Moromisato, S. Reucroft, E. von Goeler, T. Yasuda
Northern Illinois University: J. Green, D. Hedin, V. Sirotenko, S. Willis
University of Notre Dame: J. Bishop, N. Biswas, N. Cason, J. Godfrey, V. P. Kenney, J. Piekarz, R. Ruchti, W. Shepard
Oak Ridge National Laboratory: G. T. Alley, F. S. Alsmiller, R. G. Alsmiller, Jr., C.L. Britton, M. S. Emery, B.J. Frame C. Y. Fu, T.A. Gabriel, (physicist) C. W. Glover, (physicist) M.P. Hechler, R.M. Leitch, J. A. Mahall, T. L. Ryan, J.T. Shaffer, D. H. Vandergriff
Ohio State University: B. Byslma, L. S. Durkin, J. Fitch, T. Y. Ling, S. K. Park, T. A. Romanowski, C.J. Rush
Okayama University: N. Tamura
University of Oklahoma: P. Gutierrez, D. Kaplan, P. Skubic
Osaka City University: T. Okusawa, T. Takahashi, Y. Teramoto, T. Yoshida
Osaka University: J. Haba, Y. Nagashima, T. Yamanaka
Centre for Research in Particle Physics, Ottawa: M. Dixit, M. Losty, H. Mes, G. Oakham
University of Oxford: J. Bibby, R. J. Cashmore, N. Harnew, R. Nickerson, W. Williams
University of Pavia: G. Boca, M. Cambiaghi, G. Introzzi, G. Liguori, S. Ratti, P. Torre
University of Pennsylvania: L. Gladney, R. J. Hollebeek, M. Newcomer, R. Van Berg, H. H. Williams
Pennsylvania State University: T. A. Armstrong, K. W. Hartman, A. Hasan, S. F. Heppelmann, R. A. Lewis, J. N. Lim, E. D. Minor, B. Y. Oh, G. A. Smith, W. S. Toothacker, J. Whitmore, Y. Zhang
University of Pisa: S. R. Amendolia, F. Bedeschi, G. Bellittini, V. Bolognesi, M. Cobal, H. Grassmann, S. Leone, M. Mangano, A. Menzione, G. Pauletta, H. Wenzel, F. Zetti

University of Pittsburgh: E. E. Engels, Jr., T. Humanic, P. F. Shepard

Institute of High Energy Physics, Protvino: Y. Antipov, A. Denisov, A. Derevshchikov, V. Dorofeev, A. Gorin, M. Kirsanov, E. Kistenev, V. Kochetkov, I. Kotov, A. Kozelov, A. Kozshevnikov, V. Kubarovsky, A. Kulik, N. Kulman, A. Kulyavtzev, V. Kurshetsov, L. Landsberg, V. Lapshin, I. Manuilov, Yu. Mikhailov, V. Molchanov, V. Mukhin, Yu. Pischalnikov, Yu. Protopopov, A. Ronzhin, A. Ryazanzev, V. Rykalin, V. Semenov, Yu. Shamshin, A. Soldatov, V. Solyanik, N. Tyurin, D. Vavilov, V. Victorov

Purdue University: B. Abbott, V. E. Barnes, A. Bujok, D.D. Carmony, R. Davies, R. DeBonte, A. F. Garfinkel, L.J. Gutay, D. S. Koltick, A. T. Laasanen, M.C. Maher, R. McIlwain, D. H. Miller, I. Rickey, J. Ross, J. Schmitz, E. Shibata, I. P. Shipsey, B.C. Stringfellow, C. Talmadge, J. Tonnison

Rice University: D. Adams, S. Ahmad, B. Bonner, M. Corcoran, H. Miettinen, G. Mutchler, J. Roberts, J. Skeens

LAFEX/CPBF, Rio de Janeiro G. Alves, I. Roditi, A. Santoro, R. Shellard, M. Souza

University of Rochester: A. Bodek, P. De Barbaro, S. Kanda, F. Lobkowicz, W. Sakumoto, A. Sill, P. Slattery

Rockefeller University: G. Appolinari, A. Bhatti, L. Demortier, N. Giokaris, D. Goulianos, A. Maghakian, P. Melese, R. Rusack, A. Titov, A. Vacchi, Q. Wang, S. White, Z. Wu

Rutgers University: T. Devlin, T. Watts

Rutherford Appleton Laboratory: M. Edwards, S. P. H. Quinton, P. Seller

CE Saclay, France: E. Beuville, J. C. Brisson, P. Bonamy, J. Bystricky, L. Chevalier, G. Cohen-Tannoudji, G. Comby, J. Ernwein, O. Gachelin, J. F. Glicenstein, R. Hubbard, A. Joudon, P. Le Du, F. Lugiez, M. Mur, J. Poinsignon, S.P.H. Quinton, P. Rebougeard, F. Rondeaux, M. Rouger, P. Seller, B. Thooris, G. Villet

Saga University: A. Murakami, S. Kobayashi

Saitama College of Health: K. Masuda

Sofia State University, Bulgaria: R. V. Tsenov, A. B. Iordanov

Superconducting Super Collider Laboratory: D. Bintinger, F. Bird, C. Blocker, D. Coupal, D. Etherton, S. Frederiksen, A. Fry, T. Fukui, M. Hechler, K. Hess, H. Johnstad V. Kopytoff, J. Krebs, R. Lavelle, G. Leskins, T. Pal, J. Petry, J. Piles, J. Siegrist, J. Thunborg, T. Thurston, J. Tseng, M. Turcotte, N. Wells, J. Western, T. Winch, J. Yarba

Institute of Nuclear Physics, Tashkent: S. Aliev, S. Kan, A. Khaneles, A. Pak, E. Surin, B. Yuldashev

Physical Technical Institute, Tashkent: M. Alimov, K. Gulamov, V. Kaprior, V. Myalkovski, K. Turdaliev, A. Yuldashev

Institute of High Energy Physics, Tbilisi State University: N. S. Amaglobely, B. G. Chiladze, D. I. Hubua, R. G. Salukvade

Tel Aviv University: J. Grunhaus, R. Heifetz, A. Levy

University of Tennessee: T. Handler G.T. Condo

Texas A&M University: E. Barasch, A. David, H. Demroff, T. Elliott, M. Gui, F. R. Huson, T. Kamon, T. Magandar, P. M. McIntyre, Y. Pang, V. Scarpine, J.D. Smith, H.J. Trost, J. T. White

University of Texas at Arlington: P. Draper, S. P. Rosen, L. Sawyer, A. P. White

University of Texas at Dallas: C. D. Cantrell, R. C. Chaney, E. J. Fenyves, H. Hammack, J. Izen, N. P. Johnson, W. B. Lowery, J. Orgeron

Tohoku Gakuin University: M. Higuchi, Y. Hoshi

Tohoku University: K. Abe, K. Hasegawa, H. Yuta

University of Tokyo (Institute for Nuclear Study): S. Homma, S. Kato, T. Miyachi, K. Nishikawa, S. Sugimoto

Tokyo Institute of Technology: K. Kaneyuki, T. Tanimori, Y. Watanabe

Tokyo Metropolitan University: M. Chiba, R. Hamatsu, T. Hirose, S. Kitamura

Tokyo University of Agriculture and Technology: T. Emura, K. Takahashi

University of Toronto: D. C. Bailey, A. Holscher, G. J. Luste, J. F. Martin, R. S. Orr, J. D. Prentice, P. Sinervo, T-S. Yoon

Triumf Laboratory: R. Henderson, M. Salomon

University of Tsukuba (Institute of Physics): T. Asakawa, K. Hara, H. Hirade, T. Inuma, T. Kaneko, S. Kim, K. Kondo, Y. Miyamoto, S. Miyashita, I. Nakano, M. Ninomiya, S. Ogawa, M. Okabe, J. Suzuki, M. Takano, K. Takikawa, K. Yasuoka

University of Tsukuba (Institute of Applied Physics): Y. Asano, S. Mori, Y. Takada

Tufts University: T. Kafka, W. A. Mann, R. H. Milburn, A. Napier, K. Sliwa

Virginia Polytechnic Institute and State University: J.A. Kochocki, B. Lu, L. W. Mo, T. A. Nunamaker

Wakayama Medical College: M. Daigo

University of Washington: R. J. Davisson, G. Liang, H. J. Lubatti, R. J. Wilkes, T. Zhao

University of Wisconsin: J. Bellinger, D. Carlsmith, J. Cherwinka, S. Dasu, F. Feyzi, C. Fordham, C. Foudas, T. Gorski, A. Jaske, J. Lackey, R. Loveless, S. Lusin, J. Nachtman, G. Ott, D. Panescu, D. D. Reeder, W. Smith, W. Temple

Yale University: P. Cushman, S.R. Hou, V. Singh, J. Slaughter

Yasuda Woman's College: Y. Chiba

York University: W. R. Frisken, D. Hasell, R. Koniuk

Other collaborators involved in the SDC proposal preparation:

Draper Laboratory: E. Berk, J. Garman, J. Govignon, R. Martorana, M.E. Womble

IBM FSD Team:

Dallas Marketing Center: W. Courtney, S. Fisher

Houston Laboratory: A. Elam, G. Kubena, K. Liao, L. Roberts, R. Ennis

Space Sciences Laboratory, University of California at Berkeley: J. F. Arens, J. G. Jernigan

Hughes Electro-optical & Data Systems Group and Hughes Technology Center: G. Atlas, O. Barkan, T. Collins, G. Kramer, C. Pfeiffer, B. Wheeler, D. Wolfe, S. Worley

Martin Marietta Strategic Systems: T. Adams, J. Bakken, J. Brogan, B. G. Guenterberg, M. Montgomery, J. Pohlen, W. C. Vinnedge

Quantum Research Services, Inc.: W. L. Dunn, S. K. Myers, F. O'Foghludha, A. M. Yacout

Rockwell International Corporation: E. J. Anderson, M. D. Petroff

Kaiser Engineering: J. Brown, L. Dittert, W. McGinley, A. Nunez, M. Riddle

Silicon Dynamics Inc.: D. Klokow, L. VanderHave

Stanford Linear Accelerator Center: S. L. Shapiro

Westinghouse Science and Technology Center: M. A. Burke, C. W. Einolf, D. T. Hackworth, D. Marschik, D. W. Scherbarth, R. L. Swensrud, J. M. Toms

Table of Contents

1. Introduction	1-1
1.1. Brief history	1-1
1.2. Motivation	1-1
1.3. Technological choices	1-2
1.4. The collaboration	1-2
1.5. Summary	1-3
2. Summary and overview of the detector	2-1
2.1. Summary of detector parameters	2-1
2.2. Tracking system	2-3
2.2.1. Silicon tracker	2-3
2.2.2. Straw-tube barrel tracker—baseline	2-3
2.2.3. Gas microstrip intermediate tracker—baseline	2-9
2.2.4. Barrel tracker—scintillating fiber option	2-9
2.3. Superconducting solenoid	2-10
2.4. Calorimetry	2-10
2.4.1. Central calorimetry	2-11
2.4.2. Forward calorimetry	2-11
2.5. Muon System	2-13
2.6. Electronics systems and online computing	2-16
2.7. Offline computing	2-16
2.8. Test beams	2-17
2.9. Civil construction and installation	2-18
2.10. Examples of physics performance	2-18
3. Physics and detector requirements	3-1
3.1. Overview	3-1
3.1.1. Detector models	3-1
3.2. Electroweak symmetry breaking	3-5
3.2.1. Low-mass Higgs	3-7
3.2.2. Intermediate-mass Higgs	3-17
3.2.3. Heavy Higgs	3-22
3.2.4. SUSY extensions	3-35
3.2.5. Gauge boson pairs	3-37
3.2.6. Strong breaking	3-38
3.3. Physics of the t Quark	3-40
3.3.1. Detection and mass measurement for the t quark	3-42
3.3.2. Decays of $t\bar{t}$ to <i>lepton + jets</i> final states	3-45
3.3.3. Decays of the t quark to charged Higgs bosons	3-48
3.4. SUSY searches	3-59
3.4.1. Supersymmetric phenomenology	3-59
3.4.2. The missing- E_t signature	3-62
3.4.3. Squarks and missing- E_t	3-64
3.4.4. Missing- E_t and the forward calorimeter	3-65
3.4.5. Same-sign lepton signature for gluinos and squarks	3-67
3.4.6. Direct production of charginos and neutralinos	3-68
3.4.7. Conclusions	3-70
3.5. Heavy boson searches	3-70
3.6. Compositeness	3-73
3.7. QCD tests	3-75
3.7.1. Jet production	3-75

3.7.2.	Heavy quark production	3-76
3.7.3.	Prompt photon production	3-80
3.7.4.	W and Z production	3-81
3.7.5.	Multiple boson production	3-83
4.	Central tracking system	4-1
4.1.	Introduction	4-1
4.1.1.	Summary of physics requirements for tracking	4-1
4.1.2.	Summary of baseline detector configuration	4-3
4.1.3.	Fiber tracker option	4-9
4.2.	Rationale for design	4-9
4.2.1.	Matching of detector to physics requirements	4-10
4.2.2.	Material in the tracking volume and electron identification	4-18
4.2.3.	First level trigger	4-19
4.2.4.	Second level trigger	4-22
4.2.5.	Alignment requirements	4-22
4.2.6.	Possible upgrades	4-24
4.3.	Radiation environment	4-24
4.3.1.	Fluence of particles from interactions	4-24
4.3.2.	Albedo neutrons	4-25
4.3.3.	Backsplash from magnet coil	4-26
4.3.4.	Detector lifetime	4-26
4.3.5.	Electronics lifetime	4-28
4.4.	Design of detector elements	4-29
4.4.1.	Silicon detector and module	4-29
4.4.2.	Pixel detector option	4-34
4.4.3.	OTD design and modularity	4-37
4.4.4.	ITD design and module	4-41
4.5.	Electronics	4-43
4.5.1.	Silicon front-end electronics	4-43
4.5.2.	ITD electronics	4-44
4.5.3.	Straw tube electronics	4-46
4.6.	Mechanical systems and utilities	4-48
4.6.1.	STS mechanical design	4-48
4.6.2.	Silicon electronics cooling	4-59
4.6.3.	Silicon alignment strategy	4-61
4.6.4.	OTD mechanical design	4-64
4.6.5.	OTD alignment strategy	4-69
4.6.6.	ITD mechanical design	4-70
4.6.7.	ITD alignment	4-72
4.6.8.	OTD and STS supports	4-72
4.6.9.	Servicing and utilities	4-74
4.7.	Research and development plan	4-77
4.7.1.	STS R&D and program plan	4-77
4.7.2.	OTD R&D and program plan	4-79
4.7.3.	ITD R&D and program plan	4-80
4.8.	Scintillating fiber charged particle tracker	4-82
4.8.1.	Introduction	4-82
4.8.2.	Mechanical structure	4-83
4.8.3.	Material in the tracking volume	4-86
4.8.4.	System performance	4-88
4.8.5.	Development of fiber tracking technology	4-90

4.8.6.	R&D schedule	4-96
5.	Superconducting magnet	5-1
5.1.	Introduction	5-1
5.2.	Magnetic field design	5-1
5.3.	Solenoid design	5-2
5.3.1.	General design	5-2
5.3.2.	Coil design	5-6
5.3.3.	Superconductor design and development	5-10
5.3.4.	Radiation shield	5-12
5.3.5.	Cryostat vacuum vessel	5-12
5.3.6.	Coil support	5-13
5.3.7.	Chimney	5-15
5.3.8.	Thermal design and cooldown characteristics	5-16
5.3.9.	Diagnostics and instrumentation	5-17
5.3.10.	Magnet assembly and initial test	5-18
5.4.	Transportation and installation	5-18
5.5.	Cryogenics	5-18
5.5.1.	General requirements	5-18
5.5.2.	SDC cryogenic system	5-19
5.6.	Electrical system	5-20
5.6.1.	General features	5-20
5.6.2.	Power supply	5-20
5.6.3.	Discharge resistors and switches	5-21
5.7.	Prototype R&D	5-22
5.7.1.	Prototype coil	5-22
5.7.2.	Prototype vacuum chamber	5-23
5.8.	Safety and protection	5-23
5.9.	Schedule	5-23
6.	Calorimetry	6-1
6.1.	Introduction	6-1
6.2.	Physics requirements	6-2
6.2.1.	Calorimeter environment	6-2
6.2.2.	General calorimeter requirements	6-2
6.2.3.	Requirements for the shower maximum detector	6-6
6.2.4.	Requirements on the forward calorimeter	6-8
6.2.5.	Requirements summary	6-8
6.3.	Calorimeter design	6-8
6.3.1.	The central calorimeter	6-10
6.3.2.	Barrel electromagnetic calorimeter	6-11
6.3.3.	EM calorimeter constant term	6-13
6.3.4.	Endcap electromagnetic calorimeter	6-14
6.3.5.	Barrel and endcap hadron calorimeter	6-16
6.3.6.	Forward calorimeter (FCal)	6-21
6.4.	Tile/fiber optical system design	6-22
6.4.1.	Light yield	6-23
6.4.2.	Longitudinal light yield variation	6-24
6.4.3.	Transverse light yield variation	6-26
6.4.4.	Optical system fabrication studies	6-26
6.4.5.	Photon detector and associated electronics	6-31
6.5.	Radiation hardness of tile/fiber calorimetry	6-31
6.5.1.	Tests on components	6-33

6.5.2.	Radiation Hardness Tests with assembled EMC modules	6-34
6.5.3.	Radiation-hard scintillator development	6-38
6.6.	Mechanical and integration issues	6-40
6.6.1.	Introduction	6-40
6.6.2.	Barrel calorimeter	6-40
6.6.3.	Endcap calorimeter	6-42
6.6.4.	Removable hadronic plug	6-43
6.6.5.	Barrel support system	6-44
6.6.6.	Endcap support system	6-46
6.6.7.	Forward calorimeter integration and SSC interface aspects	6-46
6.7.	Shower-maximum detector and electron identification	6-46
6.7.1.	Electron and photon identification	6-46
6.7.2.	Requirements for the SMD	6-48
6.7.3.	SMD beam test results	6-51
6.7.4.	SMD design	6-55
6.7.5.	Photodetectors for SMD	6-58
6.7.6.	Calibration of the shower maximum detector	6-60
6.8.	Calibration	6-60
6.8.1.	Introduction	6-60
6.8.2.	Source calibrations	6-61
6.8.3.	Optical flasher calibration	6-61
6.8.4.	<i>In situ</i> calibration	6-62
6.9.	Engineering and R&D plan for the central calorimeter	6-62
6.9.1.	Introduction	6-62
6.9.2.	Engineering tasks	6-63
6.9.3.	Safety analysis	6-65
6.9.4.	Cost estimate development	6-65
6.9.5.	Prototypes	6-66
6.9.6.	Manufacturing plan	6-67
6.10.	R&D and engineering for the forward calorimeter	6-67
6.10.1.	High-pressure gas ionization option	6-68
6.10.2.	Liquid scintillator forward calorimeter option	6-69
6.10.3.	Milestones for 1992 and 1993; the next decision level	6-71
6.11.	Test beam results (FNAL T841)	6-71
6.11.1.	Beam test of reconfigurable-stack calorimeter	6-71
6.11.2.	The MP program	6-72
7.	Muon system	7-1
7.1.	Introduction and overview	7-1
7.1.1.	The components	7-1
7.1.2.	The trigger	7-3
7.1.3.	Muon identification	7-5
7.1.4.	Momentum measurement	7-6
7.2.	Design choices	7-6
7.2.1.	Toroid thicknesses	7-7
7.2.2.	Chamber resolution	7-8
7.2.3.	Number of layers	7-8
7.2.4.	ϕ chamber placement	7-10
7.2.5.	BW2-BW3 and FW4-FW5 separation	7-10
7.3.	Scintillation counters	7-12
7.3.1.	Central region	7-13
7.3.2.	Forward region	7-15

7.3.3.	Calibrations	7-17
7.3.4.	Upgrades	7-17
7.4.	Čerenkov counters	7-18
7.4.1.	Introduction	7-18
7.4.2.	Čerenkov counter parameters	7-18
7.4.3.	Detector simulation and trigger efficiency	7-19
7.4.4.	Mechanical structure	7-19
7.5.	Central chambers	7-20
7.5.1.	Introduction	7-20
7.5.2.	Construction philosophy	7-22
7.5.3.	Design concept	7-23
7.5.4.	Structural analysis	7-31
7.5.5.	Cell design	7-32
7.5.6.	Prototype	7-34
7.6.	The forward chambers	7-35
7.6.1.	Special background considerations	7-35
7.6.2.	Special constraints	7-36
7.6.3.	Mechanical design concept	7-37
7.6.4.	Muon drift tubes	7-38
7.6.5.	Electronics and gas deployment	7-38
7.6.6.	Alignment	7-40
7.7.	Toroidal magnets	7-40
7.7.1.	Introduction	7-40
7.7.2.	Barrel toroid	7-41
7.7.3.	Barrel coils	7-44
7.7.4.	Toroid support	7-46
7.7.5.	Forward toroids	7-46
7.7.6.	Forward coils	7-48
7.7.7.	Magnetic field analysis	7-48
7.8.	Alignment and position monitoring system	7-48
7.8.1.	Introduction	7-48
7.8.2.	System support	7-50
7.8.3.	Supermodule referencing	7-52
7.8.4.	Inter-regional referencing	7-53
7.9.	Supertower prototype	7-56
7.10.	Assembly and installation	7-56
7.10.1.	Introduction	7-56
7.10.2.	Receiving and storing	7-57
7.10.3.	Testing and maintenance	7-57
7.10.4.	Chamber and supertower assembly and alignment	7-58
7.10.5.	Supertower storage and installation	7-58
8.	Electronics	8-1
8.1.	Front-end electronics, triggering, and data acquisition	8-1
8.1.1.	Introduction to front-end electronics, triggering, and data acquisition	8-1
8.1.2.	Overview	8-2
8.1.3.	Trigger system	8-3
8.1.4.	Front-end electronics	8-4
8.1.5.	Data acquisition system	8-5
8.2.	Triggering	8-6
8.2.1.	Introduction	8-6
8.2.2.	Trigger principles	8-6

8.2.3.	Trigger structure	8-7
8.2.4.	Level 1 trigger function	8-9
8.2.5.	Level 2 Trigger function	8-25
8.2.6.	Level 3 trigger function	8-29
8.2.7.	Schedule	8-31
8.3.	Calorimetry front-end electronics	8-31
8.3.1.	Introduction/requirements	8-31
8.3.2.	Switched Capacitor Arrays readout system	8-34
8.3.3.	Digital Phototube Readout system	8-42
8.3.4.	Shower maximum detector front-end electronics	8-50
8.4.	Straw tube front-end electronics	8-54
8.4.1.	Design requirements	8-54
8.4.2.	Overview	8-54
8.4.3.	Data flow and readout	8-56
8.4.4.	Analog signal processing	8-57
8.4.5.	Bipolar preamplifier and shaper	8-58
8.4.6.	Bipolar discriminator	8-59
8.4.7.	Drift time measurement	8-60
8.4.8.	TMC	8-60
8.4.9.	Level 2 buffer	8-61
8.4.10.	Electrical interference	8-62
8.5.	Muon front-end electronics	8-64
8.5.1.	Front-end electronics location	8-64
8.5.2.	Wire chamber front end	8-65
8.5.3.	The phototube front end	8-66
8.5.4.	Data storage and data paths	8-66
8.5.5.	Trigger overview	8-67
8.5.6.	Trigger implementation	8-67
8.6.	Scintillating fiber option front-end electronics	8-69
8.6.1.	Introduction	8-69
8.6.2.	Physical overview	8-69
8.6.3.	Amplifier-shaper-discriminators	8-70
8.6.4.	Trigger and Readout board	8-71
8.6.5.	Interfaces to trigger and DAQ	8-74
8.6.6.	Readout of fibers not involved with triggering	8-75
8.7.	Intermediate tracking detector electronics	8-75
8.7.1.	Detector and signal	8-75
8.7.2.	Overview of on-detector electronics	8-77
8.7.3.	Front end layout	8-78
8.7.4.	Front-end amplifiers	8-78
8.7.5.	D.C. voltage supplies	8-79
8.7.6.	Cooling	8-79
8.7.7.	Data transmission	8-79
8.7.8.	Trigger primitives	8-80
8.7.9.	Pipelines	8-83
8.7.10.	Status	8-83
8.8.	Silicon tracker off-detector electronics	8-83
8.8.1.	Triggering	8-86
8.9.	Data acquisition system	8-88
8.9.1.	Data acquisition system requirements	8-89
8.9.2.	Overview of the data acquisition system architecture	8-94

8.9.3.	Partitioning and stand-alone operations	8-96
8.9.4.	Components of the data acquisition system	8-97
8.9.5.	System software	8-110
8.9.6.	Event data flow, its control and trigger inhibiting	8-112
8.9.7.	Error detection and recovery	8-113
8.9.8.	Stand-alone data acquisition systems	8-114
8.9.9.	Quantities and locations of data acquisition system components	8-115
8.9.10	Commissioning The Data Acquisition System (Installation, Integration and Testing)	8-116
8.9.11.	System simulations	8-116
8.9.12.	Milestones and critical path items	8-117
8.10.	Ancillary systems controls	8-118
8.10.1.	Architecture	8-118
8.10.2.	Software	8-119
8.10.3.	Subsystem special requirements	8-120
9.	Online computing	9-1
9.1.	Overview	9-1
9.2.	Functional requirements	9-1
9.2.1.	Acquiring and recording data	9-1
9.2.2.	Insuring data integrity	9-2
9.3.	Implementation	9-2
9.3.1.	Methodology	9-2
9.3.2.	Functional implementation	9-3
9.3.3.	Hardware	9-4
9.3.4.	Run control	9-4
9.3.5.	Configuration control and monitor	9-5
9.3.6.	Data logging	9-5
9.3.7.	Data monitoring	9-6
9.3.8.	Condition monitoring	9-6
9.3.9.	Calibration	9-6
9.3.10.	Physics data processing	9-6
9.3.11.	Database	9-7
9.3.12.	Documentation	9-7
9.3.13.	Software tools and standards	9-7
9.4.	Relationships and impact on other systems	9-7
9.4.1.	Front-end electronics	9-7
9.4.2.	Level 1 and 2 trigger systems	9-7
9.4.3.	Data acquisition and Level 3	9-8
9.4.4.	Offline computing	9-8
9.4.5.	Test-beam systems	9-8
9.4.6.	Ancillary controls	9-8
9.4.7.	User code	9-8
9.5.	Schedule	9-8
9.5.1.	Test beam testing	9-8
9.5.2.	Data format	9-9
9.5.3.	Sub-detector installation	9-9
9.5.4.	Level 3, data routing and logging	9-9
9.5.5.	Cosmic-ray running March 1999	9-9
9.5.6.	Detector turn-on October 1999	9-9
9.6.	Organization	9-9
9.6.1.	Project management	9-9

10. Offline computing	10-1
10.1. Introduction	10-1
10.2. Functional requirements	10-3
10.2.1. Data production requirements	10-5
10.2.2. Data storage management	10-5
10.2.3. Data analysis requirements	10-7
10.2.4. Simulation requirements	10-9
10.2.5. Communications requirements	10-9
10.3. Software requirements	10-10
10.3.1. Software environments	10-10
10.3.2. Software architectural requirements	10-11
10.3.3. Data model	10-12
10.3.4. Software engineering	10-13
10.3.5. Tools	10-13
10.4. Hardware design	10-14
10.4.1. Overview	10-14
10.4.2. Generic SDC computing model	10-14
10.4.3. Production reconstruction system	10-15
10.4.4. Simulation system	10-19
10.4.5. Physics analysis system	10-19
10.4.6. Communications	10-19
10.4.7. Hardware design verification	10-20
10.5. Software design	10-21
10.5.1. Architecture	10-21
10.5.2. Development process	10-24
10.5.3. Software sources	10-26
10.5.4. Software development process	10-26
10.6. Detector simulation	10-28
10.6.1. SDCSIM goals	10-28
10.6.2. Software design of SDCSIM	10-29
10.6.3. Current status	10-29
10.6.4. Future plans for simulation programs	10-29
10.7. Organization	10-30
10.7.1. SDC computing organization	10-30
10.8. Technology forecast	10-31
10.9. Schedule and milestones	10-31
10.10. Manpower and cost	10-33
11. Safety	11-1
11.1. Introduction	11-1
11.1.1. Overview	11-1
11.1.2. Methodology and scope	11-1
11.2. Tracking	11-2
11.2.1. Overview	11-2
11.2.2. Inner tracking	11-2
11.2.3. Outer tracking	11-3
11.3. Superconducting solenoid	11-5
11.3.1. Coil	11-5
11.3.2. Cryogenics	11-6
11.4. Calorimeter System	11-6
11.4.1. Central and endcap calorimeters	11-7
11.4.2. Forward calorimeter	11-7

11.5. Muon system	11-8
11.5.1. Drift chambers	11-8
11.5.2. Scintillation counters	11-9
11.5.3. Magnetizing coils	11-9
11.6. Electronics	11-10
11.6.1. Overcurrent protection	11-10
11.6.2. High voltage	11-10
11.6.3. Rack protection	11-11
11.7. Fire protection	11-12
11.7.1. Overview	11-12
11.7.2. Detection	11-13
11.7.3. Suppression	11-13
11.7.4. Prevention and loss control	11-14
11.8. Structures	11-14
11.9. Radioactive materials	11-15
11.10. Atmosphere management system	11-16
11.10.1. Atmosphere management requirements	11-16
11.10.2. System features	11-17
11.11. Confined spaces	11-17
11.11.1. Confined space definition	11-18
11.11.2. Inner detector access spaces	11-18
11.11.3. Detector deep pit area	11-19
11.12. Summary	11-19
12. Experimental facilities	12-1
12.1. Site layout	12-1
12.2. Surface facilities	12-1
12.2.1. Assembly building	12-1
12.2.2. Detector operations building	12-4
12.2.3. Utility building	12-6
12.2.4. Installation headhouse and installation shafts	12-8
12.2.5. Personnel & equipment access building	12-8
12.2.6. Office facilities	12-8
12.2.7. Cranes and lifting equipment	12-8
12.2.8. Temporary staging areas	12-8
12.3. Underground facilities	12-10
12.4. Services	12-10
12.4.1. Power	12-14
12.4.2. Communications	12-15
12.4.3. Water systems	12-15
12.4.4. HVAC systems	12-15
12.4.5. Gas system	12-16
12.4.6. Cryogenics	12-16
13. Installation	13-1
13.1. Overview	13-1
13.2. Installation sequence	13-1
13.2.1. Detector support system erection	13-1
13.2.2. Barrel steel erection	13-4
13.2.3. Barrel steel coil installation and Phase 1 barrel muon system installation	13-7
13.2.4. Forward toroid erection	13-7
13.2.5. Barrel calorimeter lowering and installation	13-9

13.2.6. Solenoid installation	13-9
13.2.7. End cap calorimeters lowering and solenoid field mapping	13-9
13.2.8. Phase 2 barrel muon system installation	13-14
13.2.9. Central tracking system installation	13-14
13.2.10. End cap calorimeter installation	13-14
13.2.11. Forward and intermediate muon systems installation	13-14
13.2.12. Forward calorimeter installation	13-17
13.2.13. Completed detector and accelerator components installation and commissioning commissioning	13-17
13.3. Major repair or upgrade	13-17
14. Test and calibration beam plan	14-1
14.1. Introduction	14-1
14.2. Accelerator facilities	14-1
14.3. Test and calibration beam summary	14-2
14.4. Detector subsystem tests and calibration	14-2
14.4.1. Silicon tracker	14-2
14.4.2. Straw-tube tracker	14-3
14.4.3. Gas microstrip intermediate tracker	14-3
14.4.4. Scintillating fiber option	14-4
14.4.5. Central calorimetry	14-5
14.4.6. Forward calorimetry	14-5
14.4.7. Muon system	14-6
15. Cost and schedule summary	15-1
15.1. Schedule summary	15-1
15.2. Summary of cost	15-1

1. Introduction

1.1. Brief history

The Solenoidal Detector Collaboration (SDC) grew out of several initially independent U.S. efforts concentrated at LBL, ANL, and FNAL, all of them with strong university participation, and the simultaneous activities of a number of Japanese high-energy physicists who organized and participated in a series of workshops in Japan. All of these studies were aimed at the design of a solenoidal detector for doing high- p_t physics at the SSC. At a workshop held in Fermilab in September 1989, the various groups, finding much commonality in their designs, decided to combine their activities and form a single collaboration to prepare an Expression of Interest (EoI) for submission to the SSC Laboratory. A governance document was drafted, discussed and modified at the first collaboration meeting in December 1989, and ratified shortly thereafter.

The SDC submitted its EoI in May 1990, presented its design concept to the SSCL PAC in June, and responded to ensuing PAC questions in July. The SSCL responded by requesting proponents of large high- p_t detectors to combine forces where appropriate and to submit Letters of Intent (LoI's) by the end of November. The SDC submitted its LoI, and made its verbal presentation to the PAC in December. In January 1991, the SDC detector was approved to proceed to develop a full Technical Design Report.

When the SDC was initially formed in 1989, it sent out a letter to the international HEP community inviting interested collaborators to join. Since the birth of the Collaboration, a large number of new institutions from both inside and outside the U.S. have joined the SDC. The present numbers of collaborators, including physicists and engineers, are about 525 from within the U.S. and 330 from outside the U.S. These numbers have been growing continuously, and some further growth is likely.

1.2. Motivation

The term "solenoidal detector" refers to a substantial cylindrical volume, concentric with the beam, surrounded by a solenoid coil and filled with tracking detectors. This system is capable of measuring precisely the momenta of charged tracks emitted from the interaction, within the detectors' very large angular acceptance. On the outside of the solenoid is a hermetic calorimeter with fine-sampling electromagnetic sections and somewhat coarser hadronic compartments, and with a special finely segmented detector near electromagnetic shower maximum. One of our major goals is to have excellent electron identification, and precise measurement of the energies of isolated electrons and photons. To avoid degradation of these energy measurements, the solenoid coil is designed to be very "thin," and through special weighting of the signals from the first calorimeter detection layer, we expect to reduce even further any such degradation. Outside the calorimeter is an extensive muon system, including magnetized iron toroids, tracking chambers and scintillation counters, to provide muon identification, trigger capability, and in combination with the inner tracker already mentioned, excellent momentum resolution. The precision calorimetry and tracking systems extend to pseudorapidities of 2.5, and more coarse hadronic calorimetry goes out to pseudorapidities of 6, to allow detection of non-interacting neutrals through the measurement of overall missing transverse energy.

This detector design builds upon the successful CDF experience. Its aim is to enable the measurement of the largest possible number of independent quantities for each trigger event. These include e/μ identification, sign of charge and energy measurement, detection and measurement of isolated photons, measurements of jet energies and directions, identification of jets with b -hadrons, determination of charged particle multiplicities, and detection of non-interacting neutrals. It is the ability to combine all these elements of information simultaneously for a given event that gives the SDC detector its unprecedented power, both in the ability to establish (rather than just suggest) new unexpected phenomena and in the redundant identification of interesting processes predicted by present models. If past experience with collider detectors (both hadron and electron-positron) has shown anything, it has demonstrated abundantly that with multiple independent capabilities, a detector is far better than the sum of its parts or subsystems.

Since the expected ratio of interesting events to backgrounds is far smaller at the SSC than in present colliders, we shall need all the capability that we can provide.

Clearly the design choices for such a detector always represent a balance between physics needs and available resources. The twenty-fold increase in energy and nearly 1000-fold increase in luminosity relative to the present Tevatron experience introduce unprecedented demands on speed-of-response, pattern recognition capability, excellent momentum resolution, and segmentation adequate to identify fine structures. We wish to measure quarks and gluons (jets), leptons, photons, and individual hadrons. There must be superb monitoring and calibration capability to ensure proper performance of all subsystems. The HEP community cannot afford many such detectors, hence the detector must be sufficiently robust and resistant to radiation damage to promise good performance over many years. Finally the potential for luminosity increases beyond the design value of $10^{33} \text{ cm}^{-2}\text{s}^{-1}$ must be considered. The detector must be capable, with manageable modifications, to operate at higher luminosities up to $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ with sufficient functionality to attack those physics problems whose study requires the higher luminosity. We believe that the proposed detector meets all these qualifications. Furthermore, in defining the scope of this detector and doing cost/performance optimization, we need to differentiate between scope reductions which, if the need arises, can later be removed through upgrades, and those other scope reductions whose effects remain forever. Among the latter are such issues as central tracking volume, iron toroid thickness, calorimeter depth *etc.* Even though savings can be achieved through reductions in those parameters, we believe that such reductions below the levels proposed in this document would lead to unacceptable technical and performance risks. We attempt in the detailed subsystem chapters of this Report to justify these parameter choices.

1.3. Technological choices

When the SDC detector was first presented in the EoI, there were listed five potential technologies for the central calorimetry, two for the outer central tracking (where yet another choice was a hybrid of the two), *etc.* The R&D programs sponsored by the SSCL, beginning with the generic R&D and continuing with the large subsystem efforts, eventually leading to the present detector-specific R&D activities, provided much of the technical bases for making informed choices. The criteria for choosing particular technologies include feasibility, adequacy of performance, survivability, acceptable technical risk, affordable cost, and finally the strong interest of members of the SDC to build with the chosen technology.

Many of the technological choices for the SDC detector have been made, although a few still remain for which the options have been narrowed but the final choice has not been made. The decision process in most cases involved the definition of requirements for the systems in question, the preparation of Conceptual Design Reports by the proponents of the various technologies, oral presentations, recommendations by a technically well qualified ad-hoc review committee, review and recommendations by the SDC Technical Board to the Collaboration and final ratification by the SDC executive Board.

For those areas where multiple technological options still exist, continuing R&D on the several options will be required to provide an improved technical basis for a decision. The delays involved do not impact the overall detector schedule.

1.4. The collaboration

From its inception, the SDC has involved a close partnership between physicists from the U.S. and physicists from other countries. The original steering committee which wrote the draft bylaws had British, Italian, and Japanese as well as U.S. members. Over the last two years, additional groups from the U.S. and from Brazil, Canada, China, France, Israel, Italy, U.K., and also from countries in the former Soviet Union and from Eastern Europe have joined the collaboration. These groups provide essential intellectual capital, and important financial resources. Given the limited detector resources available through the SSC Project, the SDC needs to add to those resources as extensive in-kind contributions as possible from its non-U.S. members if it is to produce a detector with all the needed capabilities.

In the preparation of the EoI, LoI, and Technical Design Report, the SDC governance has been exercised through the Spokesperson, Acting Project Manager/Technical Manager, three Deputy Spokespersons, and three different bodies (Boards) with separate roles. The Institutional Board, with one representative per collaborating institution, deals with general issues of collaboration membership, and the conduct of elections for the Executive Board. The latter which presently consists of 17 elected members of the SDC deals with all issues of scientific policy, approves all important appointments to positions of responsibility, and also approves major technical decisions. Finally the Technical Board, appointed by the Spokesperson and Project Manager with the approval of the Executive Board, consists principally of the leaders of the subsystem activities as well as other experts, and recommends on all major technical and technological decisions. Clearly as the project moves to the construction phase, a new management organization must be put into place to oversee the final design and fabrication of the detector.

1.5. Summary

The last two years have seen increasingly intense efforts by the SDC to design a detector adequately matched to the immense opportunities opened up by the construction of the SSC. While its design has drawn on recent experience with the Tevatron, the large increases in both energy and luminosity require an instrument vastly more ambitious than any built in the past. By requiring excellent capabilities in tracking, calorimetry, and muon systems, the SDC believes that its proposed detector embodies maximum redundancy, an essential feature for establishing rare new phenomena in an ocean of backgrounds.

The process of establishing potential responsibilities for non-U.S. collaborators is well under way, but formal approval from the relevant funding authorities will still take some time. The specific apportionment of U.S. responsibilities among national laboratories and universities should be accomplished over the next year.

Given adequate support, the SDC is prepared to meet the schedule of collider turn-on for physics in late 1999 with a detector properly matched to the SSC opportunities and a team ready to exploit the physics.

2. Summary and overview of the detector

In this section we present an overview of the detector, summarize its major parameters, and provide examples of its physics capabilities. A few options remain open for some of the detector subsystems, and the detailed design has not yet been completed for any of the subsystems. To proceed with a comprehensive cost estimate, we have defined a preliminary baseline detector with a single choice for each subsystem option, even when alternative options are still present in our design. These choices are indicated in the sections below. Detailed descriptions of the detector subsystems, including the alternative options where appropriate and performance summaries, are given in later chapters.

2.1. Summary of detector parameters

A detailed parameters list for the preliminary baseline detector and its options is contained in Ref. 1. An isometric view of the baseline detector configuration is shown in Fig. 2-1. A sophisticated tracking system surrounds the interaction point. The tracking system consists of an inner silicon tracker and an outer tracker. Two options are presently under consideration for the outer tracker:

1. A straw-drift-tube barrel tracker covering $|\eta| < 1.8$ together with an array of gas microstrip detectors covering the region $1.8 < |\eta| < 2.8$ (baseline option) or
2. A scintillating fiber tracker option covering $|\eta| < 2.3$.

The tracking system is contained within a superconducting solenoid that provides a peak field at the interaction point of 2.0 T. The solenoid and tracking system are surrounded by hermetic calorimetry. In the central region, $|\eta| < 3$, this calorimetry consists of scintillating tile and wavelength-shifting fiber readout with lead (electromagnetic section) or iron (hadronic sections) absorbers. A fine-grained shower-maximum detector is contained in the electromagnetic section of the central calorimeter to aid in electron and photon identification. The central calorimeter is divided into a barrel section ($|\eta| < 1.4$) and two endcaps ($1.4 < |\eta| < 3$). Hermeticity is completed by forward calorimeters covering $3 < |\eta| < 6$ at both ends of the detector. High-pressure gas ionization readout or liquid scintillator in tubes are the options under consideration for this region. A large system of magnetized-iron toroids, wire chambers and scintillation counters for muon identification and momentum measurement surrounds the calorimetry. Muon triggering and identification for $|\eta| < 2.5$ are provided by this system and muon momenta are precisely determined by a combination of measurement in the central tracking system and by deflection measurements in the iron toroids. The high data rates at the SSC require a very sophisticated electronics plant. In general, front-end circuitry is located either on or very close to the active detection elements for all systems to preserve high-rate capability. Data are stored locally on the detector and then shipped via high-speed optical fiber links to the data collection point located on the surface above the interaction hall. Data are stored, discarded or transmitted in response to trigger signals from a three-level trigger system. The Level 1 system provides triggers within 4 μ s of an event, and the Level 2 system within about 50 μ s. The Level 3 trigger is formed by an extensive array of parallel processors controlled by high level software to select events for permanent storage. An online computing and control system monitors and controls the detector.

Elevation and end views of the detector are shown in Figs. 2-2 and 2-3, respectively. The detector sits in a pit in the interaction hall and is supported by a jacking system to accommodate floor movements during installation and operation. The central calorimeter is also supported by an hydraulic jacking system to allow small movements and adjustment independent of the barrel toroid. The superconducting coil and the tracking system are attached independently to the barrel calorimeter. Electronics for the tracking system and central calorimeter are located in crates on the back of the calorimeter. Access to these electronics is obtained via pathways on either side of the forward muon system. The endcap calorimeters may be retracted by about 1.2 m for access to the tracking system. The bulk of the forward muon system remains stationary and only the chambers FW1 (Fig. 2-2) move to allow the endcap to retract. Hence routine maintenance of the detector may be performed without moving the heavy components of the forward muon system and the delicate alignment of this system can be preserved. For major repairs or upgrades to the

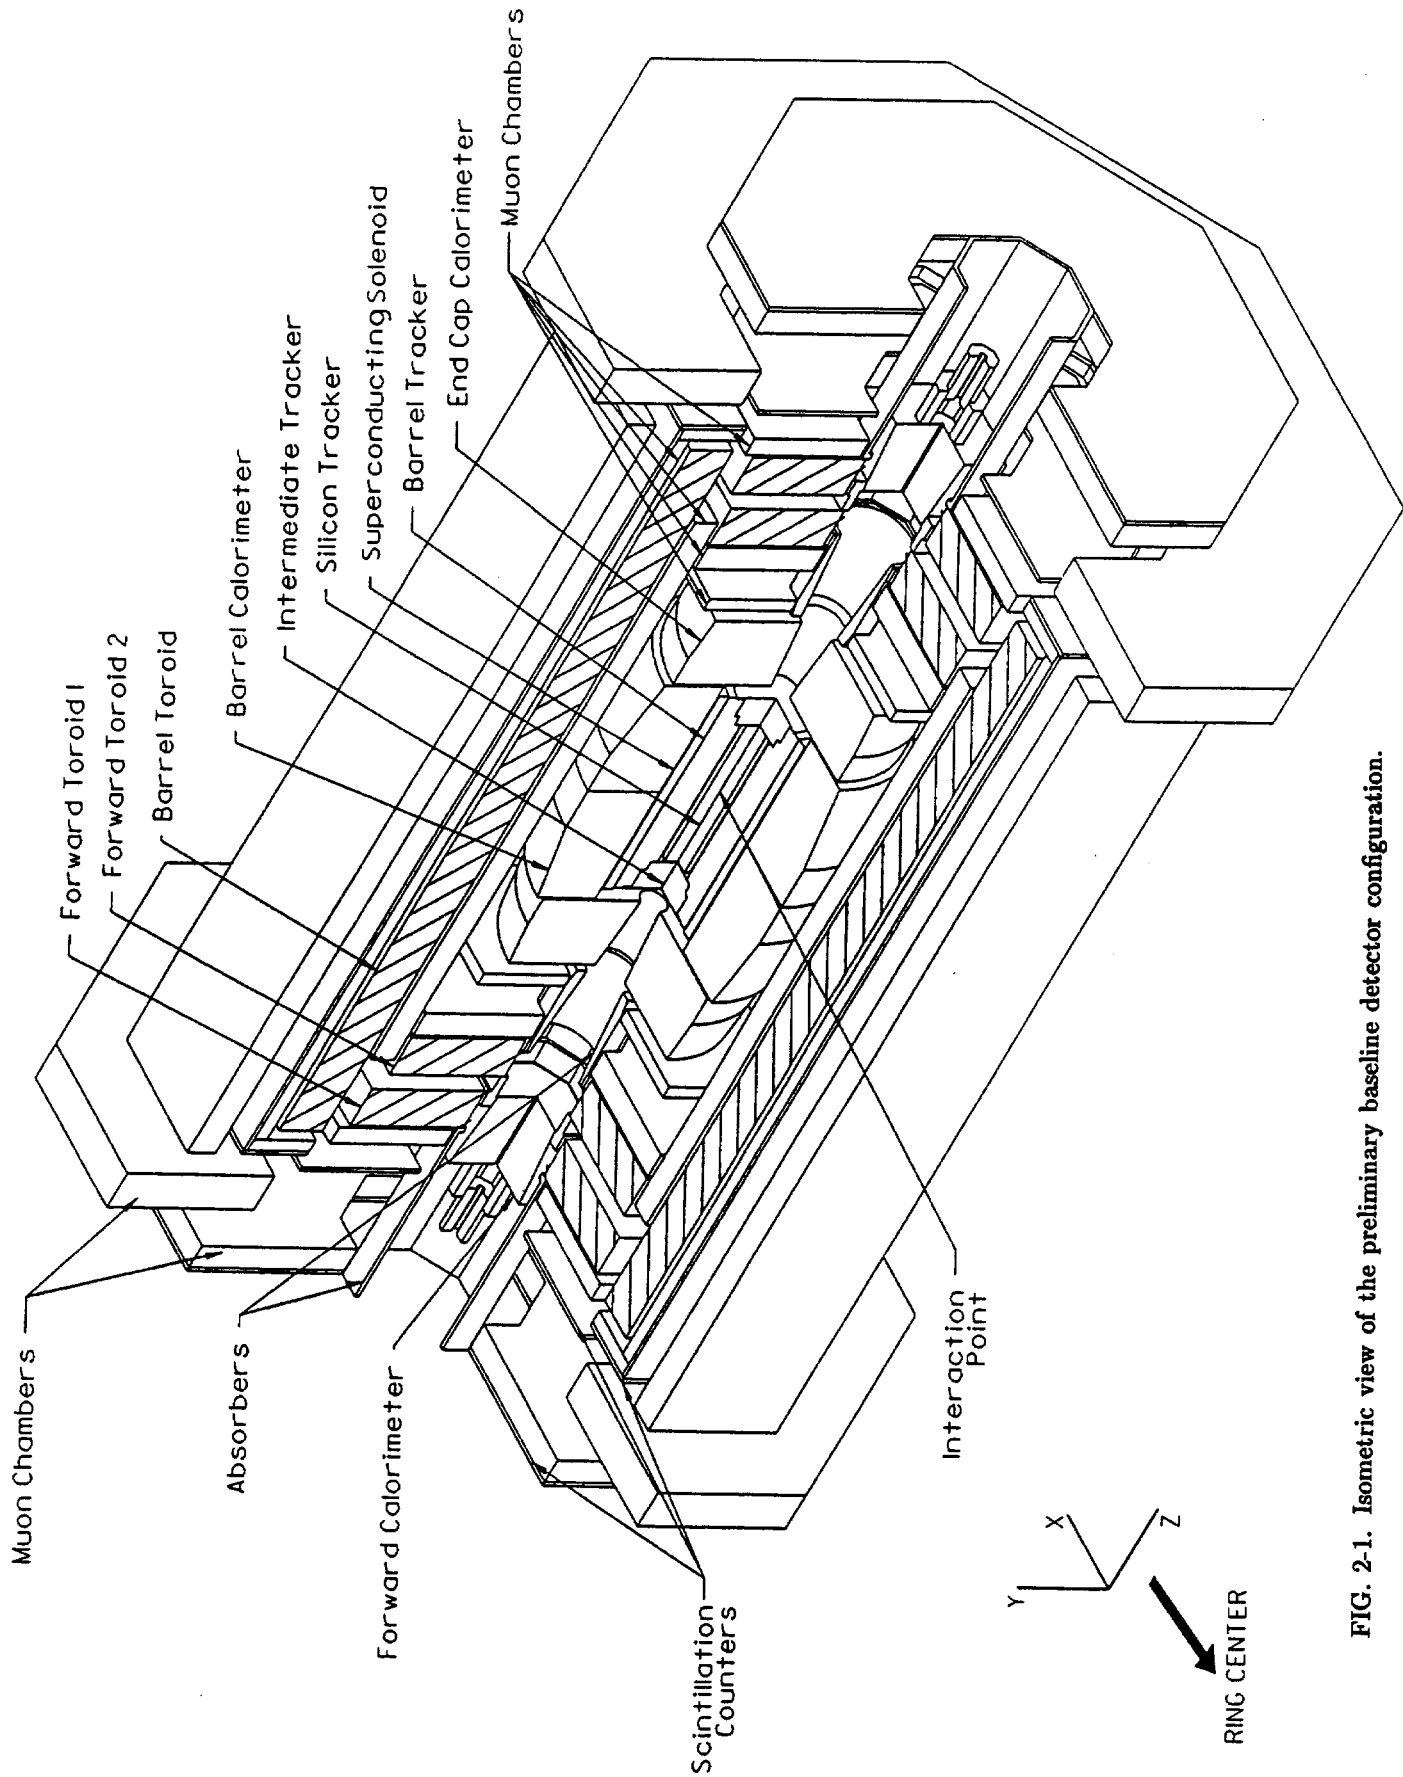


FIG. 2-1. Isometric view of the preliminary baseline detector configuration.

tracking system, the forward muon system components can be moved on rails and temporary bridges onto the operating floor of the underground hall.

A summary of the major parameters of the detector is given in Table 2-1.

2.2. Tracking system

The preliminary baseline tracking system is composed of an inner silicon tracker, a barrel straw-tube tracker and a gas microstrip intermediate tracker. A schematic drawing of the baseline configuration is given in Fig. 2-4. We are also considering the option of an all scintillating-fiber layout for the outer tracker, as shown in Fig. 2-5. The parameters of the preliminary baseline configuration and of the scintillating fiber option are summarized in Table 2-2. Each of the elements of the tracking system is described briefly below.

2.2.1. Silicon tracker

The silicon tracker consists of approximately 17 m^2 of instrumented silicon strip detectors. The silicon tracker is composed of a barrel region consisting of eight cylindrical layers of double-sided silicon strip detectors, which provide axial and small-angle stereo measurements. Thirteen double-sided disk detector arrays on each side of the barrel complete the system. Each double-sided detector is about $300 \text{ } \mu\text{m}$ thick and has a strip pitch of about $50 \text{ } \mu\text{m}$. The detectors and the on-board electronics are mounted within a low mass, highly precise space frame. This structure is in turn enclosed by a thin double-walled vessel, since cooling of the electronics heat load is provided by evaporating butane. We are considering the implementation of the two innermost layers as pixel devices as a possible option or upgrade, if research and development indicate the feasibility of doing so at a reasonable cost.

The silicon tracker covers the rapidity range $|\eta| < 2.5$, and is the key element for pattern recognition within the tracking system. In combination with the outer tracker, the track finding efficiency for isolated tracks (*e.g.*, leptons from standard Higgs particle decay) is near 100% over the entire rapidity range at design luminosity. The silicon tracker also provides the ability to reconstruct, with good efficiency, tracks of p_t above a few GeV/c even within jets of transverse energies up to a few hundred GeV . The presence of b -hadrons may be tagged with reasonable efficiency from the displaced vertices produced by their weak decays.

The silicon tracker does not contribute to the Level 1 trigger, but is used in the Level 2 trigger to provide precise information in ϕ and in p_t for combination with calorimetric information or data from the muon system.

At the design luminosity, the expected lifetime of the silicon tracker, including the on-board electronics, ranges from about 10 years at the innermost radius to about 100 years at the outermost radius.

2.2.2. Straw-tube barrel tracker—baseline

We have chosen a straw-tube barrel tracker for the preliminary baseline configuration because its performance is relatively well understood. At the present time, we believe that the combination of the established drift tube technology together with the new gas microstrip technology has lower risk than an all scintillating fiber option for the outer tracker. The status of all three technologies will be reviewed in depth in early August 1992 to determine if a final baseline design can be selected.

The layout of the straw-tube barrel tracker option has already been shown in Fig. 2-4. The 4 mm diameter straw tubes are contained in carbon-fiber-foam modules that provide a precise and rigid structure to both locate the straws and maintain the wire tension. The modules are located precisely on machined composite rings supported by carbon-fiber-foam composite cylinders. The cylinders are supported by a spaceframe composed of carbon-fiber-epoxy elements. All structural elements have been designed to be as low-mass as possible while still providing structural rigidity to maintain alignment tolerances. Both axial and stereo measurements are provided by this system. A Level 1 trigger is provided by identifying high- p_t local track vectors in any two out of the three axial superlayers. Each axial (stereo) superlayer contains




FIG. 2-2. Elevation view of the preliminary baseline detector configuration.




FIG. 2-3. End view of the preliminary baseline detector configuration.




FIG. 2-4. Schematic view of baseline tracking system configuration (silicon tracker, barrel straw-drift-tube tracker and gas microstrip intermediate tracker). Support structures and cables are not shown.

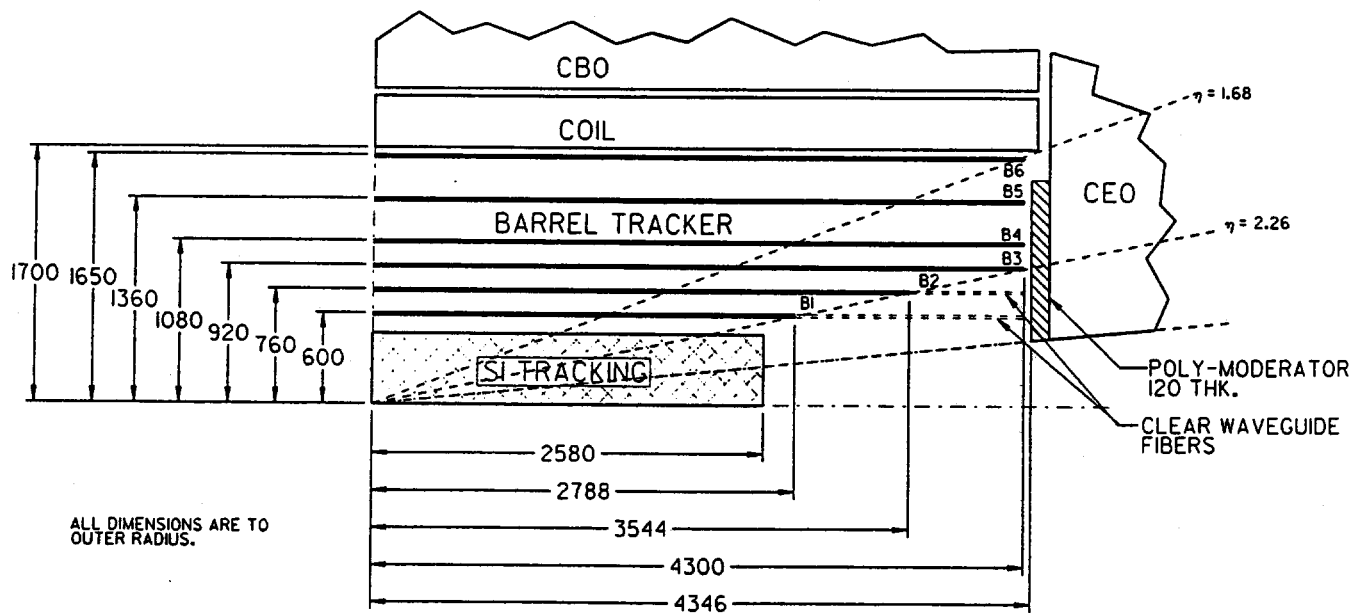


FIG. 2-5. Schematic view of the scintillating fiber option for the outer tracker. The details of the support structure are not shown.

Table 2-1
Summary of preliminary baseline detector parameters.

Tracking system	
Silicon tracker	
Number of channels	5.8×10^6
Total active area (m ²)	17
Rapidity coverage	$ \eta < 2.5$
Barrel straw-tube tracker	
Number of channels	137,164
Number of superlayers	5
Rapidity coverage	$ \eta < 1.8$
Gas microstrip intermediate tracker	
Number of channels	1.4×10^6
Number of superlayers (each end)	3
Rapidity coverage	$1.8 < \eta < 2.8$
System performance $\Delta p_t/p_t$ @1 TeV/c p_t	0.16 ($\eta = 0$), 0.60 ($\eta = 2.5$)
Calorimetry	
Barrel scintillating-tile/fiber	
Number of tower channels	10,624
Number of shower-maximum detector channels	28,672
Total depth including coil ($\eta = 0$)	10 λ
Total weight (metric tons)	$\sim 2,400$
Endcap scintillating-tile/fiber	
Number of tower channels (both ends)	9,728
Number of shower-maximum detector channels	18,432
Total weight (metric tons) (both ends)	$\sim 1,200$
Total depth ($\eta = 3$)	12.1 λ
Forward high pressure gas or liquid scintillator	
Number of tower channels (both ends)	1,056
Total weight (metric tons) (both ends)	~ 300
Total depth	12 λ
System performance	
Central calorimeter ($ \eta < 3$)	
$\Delta E_t/E_t _{EM}$	$\sim 0.14/\sqrt{E_t} \oplus 0.01$
$\Delta E_t/E_t _{HAD}$ (single π)	$\sim 0.60/\sqrt{E_t} \oplus 0.04$
Forward calorimeter ($3 < \eta < 5.5$)	
$\Delta E_t/E_t _{HAD}$	$\sim 1.0/\sqrt{E_t} \oplus 0.08$
Muon system	
Barrel toroid weight (metric tons)	16,406
Forward toroids weight (metric tons) (both ends)	4,689
Number of muon chamber channels	89,864
Number of scintillator counters	4496
System performance $\Delta p_t/p_t$ @1 TeV/c p_t	0.11 ($\eta = 0$), 0.3 ($\eta = 2.5$)
Superconducting magnet	
Field (T)	2.0
Stored energy (MJ)	146

Table 2-2

Summary of the major parameters of the preliminary baseline tracking system (silicon tracker, straw-tube tracker, and gas microstrip intermediate tracker) and of the scintillating fiber option.

Silicon tracker	
Number of channels	5.8×10^6
Number of layers (measurements on both sides)-barrel	8
Number of layers (measurements on both sides)-forward (each end)	13
Total silicon area (m^2)	17
Power dissipated (kW)	~ 7
Lifetime at design luminosity	10-100 yr (inner to outer layers)
Occupancy at design luminosity	< 0.001
Straw-tube tracker	
Straw diameter	4 mm
Number of channels	137,164
Number of axial superlayers	3
Number of stereo superlayers	2
Straw layers per superlayer	8 (axial), 6 (stereo)
Power dissipated (kW)	~ 4
Lifetime at design luminosity	> 10 yr (inner layer)
Occupancy at design luminosity	0.02-0.10 (outer to inner layers)
Gas microstrip intermediate tracker	
Number of channels	1.4×10^6
Number of superlayers (each end)	3
Measuring layers per superlayer	2 (radial) + 2 (stereo)
Power dissipated (kW) (both ends)	~ 30
Lifetime at design luminosity	To be determined by R&D
Occupancy at design luminosity	0.001-0.002
Performance of preliminary baseline option at design luminosity	
$\Delta p_t/p_t$ ($p_t = 1 \text{ TeV}/c$)	0.16 ($\eta = 0$), 0.60 ($\eta = 2.5$)
Rapidity coverage	$ \eta < 2.5$
Isolated track finding efficiency ($p_t > 10 \text{ GeV}/c$)	$\sim 99\%$
b -hadron tagging efficiency at $p_t^{\text{jet}} = 100 \text{ GeV}/c$	$\sim 30\%$
Rapidity coverage for 3σ sign determination at $p_t^{\text{track}} = 500 \text{ GeV}/c$	$ \eta < 2.5$
Scintillating fiber option	
Rapidity coverage	$ \eta < 2.3$
Fiber diameter	925 μm
Number of channels	473,200
Number of superlayers	6
Measuring layers per superlayer	4 or 8
Number of axial-only superlayers	4
Number of axial-plus-stereo superlayers	2
Lifetime at design luminosity	> 10 yr (inner layer)
Occupancy at design luminosity	0.003-0.022 (outer to inner layers)

eight (six) layers of straws. The modules are 4 m long or less with a termination, but no active electronics, at the middle at $|\eta| = 0$.

A gas mixture of tetrafluoro-methane (80%) and isobutane (20%) is used in the straws. The mixture provides a maximum drift time of about 30 ns, which is reasonably matched to the SSC interaction rate. The straw-tube cathodes are very thin copper-coated Kapton, which has been demonstrated to have better radiation resistance than aluminum cathodes. The expected lifetime of the straw-tube system, including the front-end electronics on the ends of the straws, exceeds ten years at the minimum superlayer radius at the design luminosity.

The pattern recognition and track finding capability of the straw system, in conjunction with the silicon tracker, are under study and preliminary results are reported in Chapter 4. At the design luminosity, the expected occupancy of the innermost (outermost) straw superlayer is 0.10 (0.02). Preliminary pattern recognition studies indicate that isolated tracks, for example, leptons from standard Higgs decay, can be found with efficiency $> 97\%$ by the combined silicon-straw system even at six times design luminosity.

2.2.3. Gas microstrip intermediate tracker—baseline

In the rapidity interval $1.8 < |\eta| < 2.8$, we propose to use a new technology, gas microstrip detectors. A gas microstrip detector (GMD) consists of fine metallic anode and cathode traces placed on a thin substrate (*e.g.*, glass) separated by a gap of a few millimeters from an electrode to provide a drift region. High voltage connections are made to the cathodes, drift electrode and substrate and signals are read out on the anodes. The anode pitch is typically a few hundred microns and in our design, varies with rapidity. The GMD technique has spatial resolution, two-track resolution, and speed of response that are well matched to the requirements of the intermediate tracker.

The gas microstrip intermediate tracking detector (ITD) consists of three sets of planes (superlayers) on either side of the silicon and barrel trackers. A layer in the ITD consists of an array of gas microstrip tiles approximately 15 cm on a side finely segmented in ϕ . Each superlayer consists of two layers with radial anodes (ϕ measurements) and two layers with stereo anodes of opposite inclination to provide a local space point. A Level 1 p_t -sensitive trigger is formed from the radial layers by measuring the change in ϕ from one superlayer to the next, in coarse η bins defined by the tile dimensions.

Gas microstrip detectors are a new technology that has not been used extensively in previous experiments. An aggressive international research and development program is underway to demonstrate that these detectors can be used on the scale envisioned for the ITD and that they have adequate lifetime for use at the SSC. The ITD resides in a region where the annual radiation dose is up to 10 krad at design luminosity. Research and development is required to demonstrate survivability of the GMD in such an environment.

2.2.4. Barrel tracker—scintillating fiber option

The layout of the fiber tracker option has already been shown in Fig. 2-5. Doublets of scintillating fibers are precisely arrayed on the inside and outside of carbon-fiber foam composite cylinders. The cylinders are held by a precise composite framework located at the ends of the cylinders. Both axial and small-angle-stereo measurements are provided by the fibers. The scintillating fibers are coupled to clear fibers that transmit the light to solid state photosensors, Visible Light Photon Counters (VLPC's) that are located on the outside of the central calorimeter. The VLPC's have high quantum efficiency (up to 80%) and are located in helium cryostats to maintain the 7 K temperature required for their operation. Electronics for the fiber tracker are also located on the back of the central calorimeter.

The scintillating fiber option has an occupancy that is significantly less than the straw-tube option and thereby might be expected to have better performance at luminosities greater than the design value. The trigger is implemented by correlating signals in the inner three superlayers, providing uniform trigger coverage up to $|\eta| < 2.3$. However, there is somewhat greater material, on average, in the fiber option although concentrations of material due to electronics and supports in the straw-tube/GMD option are

eliminated by the design. There is also a reduction in rapidity coverage in the present fiber design, which covers only $|\eta| < 2.3$.

A substantial R&D program is underway to demonstrate the feasibility of fiber readout with VLPC arrays on a large scale. Small scale tests have been done that demonstrate the basic concept, but large VLPC arrays with high quantum efficiency remain to be tested. Such tests are planned for 1992. The progress of this R&D and of efforts to develop accurate fiber ribbon placement will be reviewed in early August 1992, along with the progress in the straw-tube/GMD outer tracker. If sufficient data are available at that time, one option will be selected.

2.3. Superconducting solenoid

The tracking system is enclosed within a thin superconducting solenoid that provides a peak field of 2.0 Tesla. The principal parameters of the solenoid are given in Table 2-3. A research and development program is underway to verify the design of the solenoid and to make the thinnest feasible structure. This program will culminate in late 1993 with the construction of a quarter-length full-radius prototype that can be tested to simulate the compressive forces anticipated in the final coil.

Table 2-3
Parameters of the superconducting coil.

Inner radius of cryostat (mm)	1700
Outer radius of cryostat (mm)	2050
Total length of cryostat (mm)	8726
Mean conductor radius (mm)	1810
Central magnetic field (Tesla)	2.0
Nominal operating current (amps)	8000
Maximum temp. after quench ($^{\circ}\text{K}$)	< 100
Maximum voltage after quench (V)	< 500
Stored energy (MJ)	146
Thickness at $\theta = 90^{\circ}$	$1.2 X_0, 0.25 \lambda$
Total weight (tonnes)	25
Cold mass (tonnes)	20
$\int B \times dl$ (tesla-meters at 90°)	3.4

2.4. Calorimetry

The goals of the calorimeter systems are to provide electron and photon identification and energy measurement (in conjunction with the tracking system), to measure the energies and directions of jets and to provide hermetic coverage for missing transverse energy measurements. In the central region ($|\eta| < 3$) we have chosen scintillation calorimetry with lead absorber and iron absorber for the electromagnetic and hadronic sections, respectively. The scintillating detection elements are divided into tiles, each tile being read out by a waveshifting fiber. The fibers are brought to the rear of the calorimeter, bundled, masked and read out by photomultiplier tubes. In the forward region ($3 < |\eta| < 6$), we are considering two options: high pressure gas ionization readout or liquid scintillator in small tubes.

2.4.1. Central calorimetry

An elevation view of the central calorimeter is shown in Fig. 2-6 and the parameters are summarized in Table 2-4. The central calorimeter is composed of a barrel section, which in turn is built in two halves, and two endcap sections. Each endcap section has a removable "hadronic plug" covering the high rapidity region (about $2 < |\eta| < 3$) that is most susceptible to radiation damage. In the barrel region there is a single electromagnetic depth segment, which can be upgraded to two depth segments by rerouting fibers and adding phototubes. In the endcap section, there are two electromagnetic depth segments to allow for better correction of radiation-damage effects, which are more important in this region. In both the barrel and endcap, the iron hadronic absorber is segmented into two depth compartments (HAC1 and HAC2). The transverse segmentation is $\delta\eta \times \delta\phi = 0.05 \times 0.05$ in the electromagnetic sections and 0.1×0.1 in the hadronic sections, except near $|\eta| = 3$ where the granularity is coarser.

A shower maximum detector (SMD) composed of crossed strips of scintillator about 1.2 cm wide is located near the shower maximum point in both the barrel and endcap regions. The SMD aids substantially in the identification of electrons and photons by measuring the shape and location of the electromagnetic shower. The SMD is also used in the trigger to provide correlations with the central tracker.

Both the tiles in the tower segments and the strips in the SMD are read out by waveshifting fibers embedded in grooves located in each tile or strip. The fibers are routed to the back of the calorimeter. For the tower segments they are bundled and masked on a fiber-by-fiber basis using filters placed between the fiber bundle and the photomultiplier tube. This masking technique can smooth out variations from tile to tile to provide a more uniform response in depth. The required degree of masking as well as the responses of the tiles are determined by an extensive system of remotely movable radioactive sources that can illuminate and calibrate all tiles. The fibers from the SMD are read out by multianode phototubes (baseline design) or by avalanche photodiode arrays. Local electronics for the calorimeter (and the central tracker) are mounted in crates on the back of the calorimeter to minimize the high-bandwidth cable paths.

Scintillation calorimetry has the advantage of an intrinsic speed of response that comes close to matching the 16 ns time between crossings at the SSC. However, degradation of light output from radiation damage is an obvious concern. At design luminosity the maximum dose at electromagnetic shower maximum is about 6 krad/yr in the barrel region. Irradiation of electromagnetic modules in intense electron beams in China, Japan and France have demonstrated that readily available scintillators will allow electromagnetic energy measurements in the barrel region for about 100 years at design luminosity with little degradation in resolution. The radiation dose varies strongly with polar angle, approximately as $(1/\theta)^3$, and thus the annual dose in the forward part of each endcap will be large. We are continuing the development of new scintillators with increased resistance to radiation, and designing the endcap region so that the scintillator in the region of highest dose can be removed and replaced. At the design luminosity, replacement would be necessary every few years if new scintillators cannot be developed. Preliminary results on radiation resistance of small samples of newly developed scintillators indicate potential lifetime improvement factors of two or more, but large scale tests are required to validate these results.

2.4.2. Forward calorimetry

The forward calorimeter covers the rapidity range from $|\eta| = 3$ to about $|\eta| = 6$. Measurement of jet energies and angles in this region is critical to the measurement of missing transverse energy. In addition, tagging the presence of jets in this rapidity region may reduce backgrounds in the observation of signals in the central detector. The energy resolution and segmentation requirements for the forward calorimeter are not as stringent as in the central region. Table 2-5 summarizes the energy resolution and segmentation parameters of our current design. The forward calorimeter is located about 12.5 m from the interaction point. With the segmentation given in Table 2-5, the angular resolution provides adequate measurement of missing transverse energy, and identification of forward jets.

Radiation doses are much higher in the forward direction, and the feasibility of operation under such extreme conditions to a large extent determines the technologies that can be employed. Two options are




FIG. 2-6. Elevation view of the central calorimeter. The barrel calorimeter displays one line per longitudinal cell (absorber boundary), whereas the endcaps display two lines per cell (which are partly cut away).

Table 2-4
Parameters of the central calorimeter.

Barrel calorimeter	<u>EM</u>	<u>HAC1</u>	<u>HAC2</u>	
Absorber material	Pb	Fe	Fe	
Number of absorber plates	29	28	15	
Absorber thickness (mm)	4	23.95	53.90	
Number of tile layers	30	28	15	
Tile thickness (mm)	4	4	4	
Depth (at $ \eta = 0$) (not including coil)	$21X_0$ (0.85 λ)	4.14 λ	4.91 λ	
Segmentation ($\delta\eta = \delta\phi$)	0.05	0.1	0.1	
Number of channels	7186	1792	1664	
<hr/>				
Barrel modules	64			
Barrel weight (tonnes)	~ 2400			
Shower maximum detector				
Number of channels	28672			
Depth in calorimeter (not including coil)	5 X_0			
<hr/>				
Endcap calorimeter	<u>EM1</u>	<u>EM2</u>	<u>HAC1</u>	<u>HAC2</u>
Absorber material	Pb	Pb	Fe	Fe
Number of absorber plates	6	17	20	11
Absorber thickness (mm)	6.0	6.0	42.0	90.0
Number of tile layers (including SMD)	9	17	20	11
Tile thickness (mm)	4	4	4	4
Depth (at $ \eta = 3$)	$6.9 X_0(0.3\lambda)$	$18.3 X_0(0.8\lambda)$	5.04 λ	5.99 λ
Number of channels (both ends)	3200	3200	1536	1792
Segmentation ($\delta\eta = \delta\phi$) (η dependent)	0.05–0.2	0.05–0.2	0.1–0.2	0.1–0.2
<hr/>				
Endcap modules	64			
Endcap weight (tonnes) (both ends)	~ 1200			
Shower maximum detector				
Number of channels	18432			
Depth in calorimeter	7 X_0			
Hadronic endplug				
η boundary	2			
Total weight (tonnes) (both ends)	180			

under consideration for the sampling medium in the forward calorimetry: high pressure gas (about 100 atm of argon) and liquid scintillator in glass tubes. In both cases the sampling medium may require periodic replacement after accumulation of large doses of radiation.

2.5. Muon system

The muon system provides the capability to identify muons, trigger on them, and make independent measurements of muon momenta. Large magnetized-iron toroids (see Fig. 2-7) cover the rapidity range $|\eta| < 2.5$. Drift tube chambers measure the deflections of muons in the iron toroids and scintillation counters provide a precise timing signal to tag the bunch crossing of interest. At design luminosity, the primary muon momentum measurement in the central rapidity region is performed by the central tracker. In the forward region, the muon system itself has better momentum measurement capability at high p_t , since the central tracker resolution is poorer at high rapidity.

Table 2-5
Parameters of the forward calorimeter.

Technology:	High-pressure gas or Liquid scintillator
z position (mm)	12476
Inner radius (mm)	50
Outer radius (mm)	1500
Weight (both ends)	300 tonnes (approx)
Number of depth segments	2
Number of channels (both ends)	1056
Total thickness	$\geq 12 \lambda$
Segmentation ($\delta\phi = \delta\eta$)	0.2-0.8 (η dependent)
$\Delta E_t/E_t _{\text{HAD}}$	$\sim 1.0/\sqrt{E} \oplus 0.08$

The barrel iron toroid is composed of large iron segments bolted and welded together. The barrel toroid sits on a support structure that is designed to accommodate to both long-term floor motion and short-term motion from the movement of the remaining detector components into the toroid. The thickness of the toroid is the minimum depth needed to provide a reasonable Level 1 trigger rate and good muon detection efficiency.

The forward toroids are octagons with inserts to make the field as uniform as possible. Muons in the forward direction typically have higher momenta than those in the central region and greater stand-alone momentum measuring precision is also required. Three meters of magnetized iron is just sufficient to provide adequate measuring power for TeV muons.

In the baseline concept, all muon chambers consist of round drift tubes with field-shaping electrodes, which provide a near-linear time-to-distance relationship with the appropriate gas mixture (for example Argon-CO₂) and thus better spatial resolution (about 250 μm) than simple drift tubes without field shaping. In addition, the field shaping allows for two-track resolution of about 5 mm, which is needed to find muon tracks in the presence of electromagnetic debris created by the passage of multi-hundred GeV muons through the iron toroids and the chamber walls. The diameter of the drift tubes is larger in the barrel and intermediate regions, where the muon rates are lower than in the forward region. In the barrel and intermediate regions, the chamber elements are packaged as supermodules on the surface, lowered into the underground hall and mounted on the barrel toroid. A similar procedure is used for the forward system. Alignment systems are used throughout to calibrate the plane-to-plane alignment to an accuracy of about 150 μm in the barrel/intermediate region and the forward region.

Measurements in the muon chamber system are primarily for determining the muon deflection in the toroids (θ measurements), but ϕ and stereo measurements are also made in barrel/intermediate region and stereo measurements in the forward region. Stereo measurements are needed to associate tracks in the non-bend direction. The ϕ measurements are used for pattern recognition and, in association with the central tracker, to improve the momentum measurement precision at high transverse momentum.

A p_t -sensitive Level 1 trigger is formed by measuring the track deflection due to the toroids in the outer chamber layers (BW2/BW3, IW2/IW3 and FW4/FW5). The drift tubes are arranged to be projective to the interaction point. The measurement of drift-time differences between selected planes provides information related to the transverse momentum of the muon. Trigger p_t thresholds can be varied by selecting different windows in the time differences. Since the drift time in the tubes can be as much as 1 microsecond, the scintillators are used to identify the correct beam crossing. There is a single layer of scintillation counters, each with two phototubes, in the barrel/intermediate region and two layers of counters, each with one phototube, in the forward region, where rates are higher. We are considering




FIG. 2-7. Layout of the muon system.

the option of using a Cerenkov counter in the forward direction to reduce the sensitivity to neutron backgrounds, but further study is required.

The parameters of the muon system are summarized in Table 2-6.

Table 2-6
Summary of the major parameters of the muon system.

Toroids			
	<u>Barrel</u>	<u>Forward</u>	
Thickness (mm)	1500	2 × 1500	
Total weight (tonnes)	16406	4689 (both ends)	
Chambers			
	<u>Barrel</u>	<u>Intermediate</u>	<u>Forward</u>
Number of channels	44,376	13,248	32,240 (both ends)
Scintillators			
	<u>Barrel</u>	<u>Intermediate</u>	<u>Forward</u>
Number of PMT channels	3,840	640	2256 (both ends)
PMT per counter	2	2	1

2.6. Electronics systems and online computing

The general flow of data from detector elements to permanent storage is given in Fig. 2-8. Front-end electronics will be designed to match the requirements for each distinct detector subsystem. All detector subsystems require the design and fabrication of application-specific-integrated-circuits (ASIC's) to meet performance requirements for the front-end systems. Specific front-end circuits are required for the silicon tracker, the straw-tube tracker or the fiber tracker, the gas microstrip tracker, the calorimetry and the muon chambers. The circuitry for the gas microstrip tracker shares many features with the silicon design, and the muon front-end circuitry is similar to but less complex than the straw-tube circuitry.

The information from the detector subsystems that is used by the trigger system is summarized in Table 2-7. Correlations among the trigger elements at Level 1 are used to form a complex array of triggers. At Level 2, additional information and correlations are added to reduce the rate flowing into the Level 3 processor farm. Selection at Level 3 is controlled entirely by software, which may include near-complete event reconstruction. The data acquisition and online computing systems control the flow of data from the detector to the Level 3 farm and its subsequent permanent storage, and provide the interface to control all detector subsystems.

2.7. Offline computing

The principal challenges for offline computing are the storage of and access to the vast amount of data that will be accumulated during the operation of the experiment, and the management of the software development for event reconstruction and analysis by the diverse international community of the SDC. The actual processing hardware requirements of about 10^5 mips are probably not as challenging, but will represent a significant investment.

The development of code for the detector is envisioned to be divided into two broad categories. First, there will be a "kernel" of software that provides the structure to which other code may be attached. This development requires close cooperation of a core group, including software professionals, and is envisioned to occur primarily at the SSC Laboratory. The second category comprises the detector-specific codes that

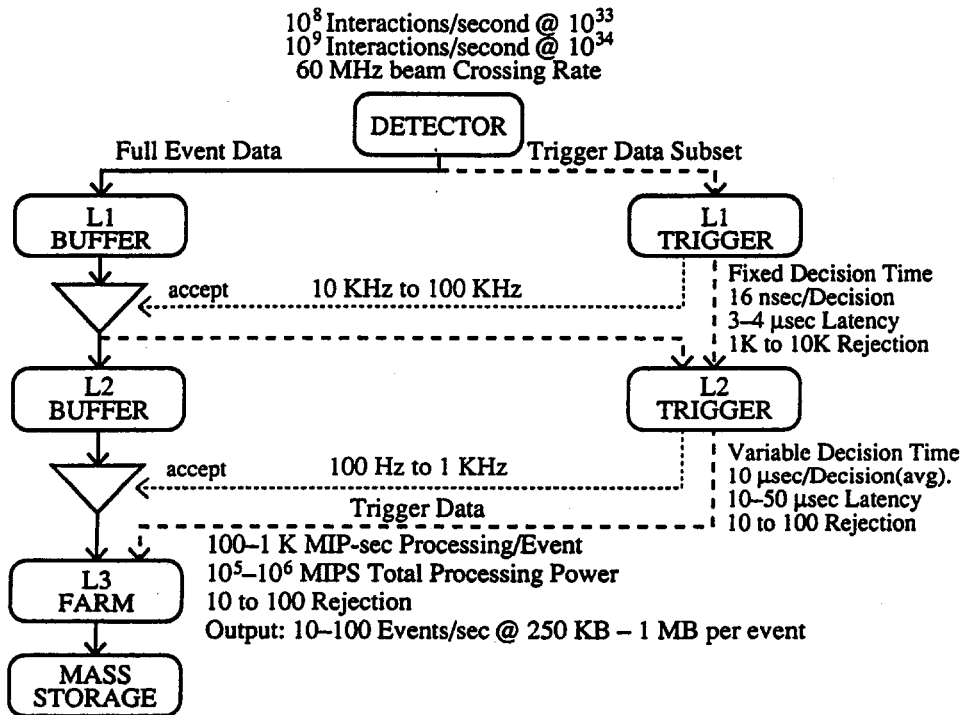


FIG. 2-8. Schematic representation of the data flow from the detector to permanent storage.

will be developed throughout the collaboration by physicists and a few software professionals within the collaboration. One must start now to develop the core group at the SSC Laboratory in order to have sufficient personnel in place to begin development of the "kernel" software in 1993.

At present, the processing hardware is envisioned to be an extension of the RISC-based multiple independent processor "farm" or "ranch" system in existence at a number of institutions, including the SSC Laboratory. The bulk, if not all, of this system would be located at the SSC Laboratory. Entry into this system would be provided, as now, via individual workstations either from local area or wide area networks.

Data storage and access to data perhaps represent the most challenging problems. Automated data storage systems containing upwards of 10^4 volumes will be required. More challenging is the development of suitable databases to allow fast access to data in response to very diverse sets of interests and criteria from the hundreds of physicists who will eventually be analyzing the experiment.

2.8. Test beams

Our plan for test-beam utilization is given in Chapter 14. The SDC plans to use test beams for radiation-damage tests, calorimeter testing and calibration, tracking system and muon system tests. Test beams in China, Japan, Europe, Russia and the United States, including low-energy electron beams for radiation-damage measurements, will be utilized during the next decade for the SDC. The major effort for testing and calibrating calorimeter modules will be done at Fermilab and eventually transferred to the test beam facilities at the SSC Laboratory medium-energy booster.

Table 2-7

Summary of detector subsystem information to be used by the Level 1 and Level 2 trigger systems.

	Level 1 Segmentation per half/end	Level 1 Data/bin	Level 2 Segmentation per half/end	Level 2 Data/bin
Silicon	—	—	2048 $\phi \times 6 \eta$	5 p_t bits
Barrel tracking	64 ϕ	2 p_t bits	1024 ϕ	4 p_t bits
Intermediate tracking	64 ϕ	1 p_t bits	1024 $\phi \times 4 \eta$	4 p_t bits
Cal. $ \eta < 1.8$				
Cal. towers	64 $\phi \times 18 \eta$	EM & HAC 8 bits Egy	64 $\phi \times 18 \eta$	EM & HAC 8 bits Egy
Shower maximum	64 $\phi \times 9 \eta$	Hit flags	1024 $\phi \times \eta$	8 bits Egy
Cal. $1.8 < \eta < 2.6$				
Cal. towers	64 $\phi \times 8 \eta$	EM & HAC 8 bits Egy	64 $\phi \times 8 \eta$	EM & HAC 8 bits Egy
Shower maximum	64 $\phi \times 4 \eta$	Hit flags	512 $\phi \times 4 \eta$	8 bits Egy
Cal. $2.6 < \eta < 3.0$				
Cal. towers	32 $\phi \times 2 \eta$	EM & HAC 8 bits Egy	32 $\phi \times 2 \eta$	EM & HAC 8 bits Egy
Shower maximum	64 $\phi \times 2 \eta$	Hit flags	256 $\phi \times 2 \eta$	8 bits Egy
Cal. $ \eta > 3.0$	8 $\phi \times 4 \eta$	EM & HAC 8 bits Egy	8 $\phi \times 4 \eta$	EM & HAC 8 bits Egy
Muon $ \eta < 1.0$	32 $\phi \times 5 \eta$	2 p_t bits	1024 $\phi \times 5 \eta$	5 p_t bits
Muon $1.0 < \eta < 2.5$	32 $\phi \times 8 \eta$	2 p_t bits	1024 $\phi \times 8 \eta$	5 p_t bits

2.9. Civil construction and installation

A summary of the status of the design of the underground hall and the surface buildings required to assemble, install and operate the detector is given in Chapter 12 and the installation plan is described in Chapter 13. Only final assembly of detector components will occur at the interaction region site. All production of calorimeter wedges, individual muon chamber components, scintillator, electronics, *etc.* will occur at the home institutions of the collaboration members. Objects too large or too heavy to transport will be assembled at the interaction region. Final assembly of tracking systems that are too delicate to transport long distances will also occur at the interaction region site.

A detailed installation plan has been devised and indicates the need for beneficial occupancy of the underground hall in January 1996. Installation and checkout of the detector components as well as the interaction region accelerator components is expected to require more than three years.

2.10. Examples of physics performance

The range of physics accessible to the SSC, with its 20-fold increase in collision energy and 1000-fold increase in luminosity over the present generation of hadron colliders, is immense. Figure 2-9 provides a visual survey of the physics processes that are described in more detail in Chapter 3. The total inelastic cross section gives an interaction rate of 10^{15} events per SSC year (defined to be 10^7 seconds of operation at the design luminosity of $10^{33} \text{cm}^{-2} \text{s}^{-1}$). Most of these events involve small momentum transfer and do not probe the mass scales of interest. Nevertheless, processes involving production of heavy objects which are rare at today's hadron colliders (the CDF detector at Fermilab recorded about 40 $W \rightarrow e\nu$ events per




FIG. 2-9. Examples of total cross sections at the SSC. The $jet-jet$ and $\gamma-\gamma$ cross sections are for wide-angle jet or γ pairs ($|\eta_{jet}|$ and $|\eta_\gamma| < 2.5$) with invariant mass greater than M . For heavy-quark pair (gluino pair) production, the cross section is evaluated at $M = M_Q$ ($M = M_{\tilde{g}}$). The Higgs cross section assumes $M_{top} = 150$ GeV. The left scale is total cross section divided by 100 mb, the approximate total pp cross section, and so the numbers are approximate production probabilities per collision. The scale on the right is the number of produced events per year under "standard conditions," defined as operation at $\mathcal{L} = 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ for 10^7 s. These rates must be further downrated by branching fractions and experimental cuts to obtain measurable rates. References and further discussion are given in Chapter 3.

day towards the end of their last run) become commonplace at the SSC (the SDC detector will be capable of recording $W \rightarrow e\nu$ events at a rate of 10 Hz at design luminosity). A second example is the production of the t quark. For $M_{\text{top}} = 150$ GeV, Fermilab would produce about 100 events during the next several years, whereas the SSC would produce 10^8 events per SSC year.

The largest interesting cross section at the SSC is that for the production of two jets. One event out of 10^4 (i.e., a rate of 10^4 Hz) has two jets with a dijet mass of greater than 400 GeV. This cross section is 10^7 times larger than that for photon pair production, serving as a reminder that robust photon and lepton identification are essential for SSC physics. Heavy quarks and other colored objects such as gluinos would be produced with large cross sections. Even for a mass of 1 TeV, there are at least 10^4 events produced per year. Heavy new Z bosons are also prolifically produced, with the observable cross section extending out to a mass of 4 TeV. Finally, the Higgs production cross section is very small. At most one event out of 10^9 would contain a Higgs boson, and the branching ratios useful for its detection are also small. For the decays of the Higgs to two photons or to four leptons, the branching ratios are typically 10^{-3} .

In the sections below, we briefly describe the physics capabilities of the proposed SDC detector, and summarize its performance in Table 2-8.

Electroweak symmetry breaking

The single most important physics issue for the SSC is the study of electroweak symmetry breaking. In the context of the Minimal Standard Model, the existence of a fundamental scalar field provides the symmetry breaking mechanism. In this case, a single Higgs boson is the only observable particle associated with the symmetry breaking sector, and its mass is the only unknown parameter. It is imperative that a general-purpose SSC detector be capable of observing such a Higgs boson at any allowable mass in order to either verify its existence, or to rule it out and force consideration of alternate mechanisms.

The search for the Standard Model Higgs divides naturally into three mass regions, each with its associated strategy. For the low mass region ($80 < M_{\text{Higgs}} < 130$ GeV), the dominant Higgs decay modes are $H \rightarrow b\bar{b}$ and $H \rightarrow \tau^+\tau^-$, which are both overwhelmed by backgrounds from the decays of t quarks. The most useful mode in this kinematic region is the rare decay $H \rightarrow \gamma\gamma$, which occurs at next-to-leading order through loop diagrams. The Higgs itself is very narrow in this region (the width is less than 100 MeV for Higgs masses below 160 GeV) so that this decay mode provides a very distinctive signature. However, the direct production of the Higgs through gluon fusion suffers from a large background of QCD continuum production of photon pairs. The production rate for a Higgs in association with a W or $t\bar{t}$ pair is suppressed by a factor of 10–20 compared to the gluon-fusion rate, but the presence of an additional high- p_t lepton from the W or t decay provides significant background suppression. A complete analysis of the associated production processes shows that the SDC detector, studying the $\ell\gamma\gamma$ final state, should be capable of discovering a Higgs in the low mass region within a single SSC year (see Fig. 3-11 and Fig. 3-13).

For the intermediate mass region ($130 < M_{\text{Higgs}} < 180$ GeV), the branching ratio for $H \rightarrow ZZ^*$ becomes significant (the $*$ denotes a virtual particle). This decay mode provides a very distinctive signature of four isolated high- p_t leptons, with little background. The SDC detector, studying this final state, should be able to observe a Higgs anywhere in the indicated mass region after one SSC year (see Fig. 3-18).

For the heavy mass region ($180 < M_{\text{Higgs}} < 800$ GeV), the WW and ZZ decay modes dominate. In the lower part of this mass range, discovery via the $H \rightarrow ZZ \rightarrow 4\ell$ mode appears straightforward (see Fig. 3-22 and Fig. 3-23). As the Higgs mass increases, the cross section for its production decreases, and its width increases dramatically (an 800 GeV Higgs has a width of 270 GeV), making discovery more difficult. We have studied the $H \rightarrow ZZ \rightarrow 4\ell$ and $H \rightarrow ZZ \rightarrow 2\ell 2\nu$ decay modes in detail. The latter has six times the event rate of the former, but requires particular scrutiny because of the requirement of observing the missing transverse energy from the neutrinos. The conclusion is that, through a combination of these two final states, a Higgs with a mass of ≤ 800 GeV should be observable within one SSC year (see Fig. 3-25 and Fig. 3-28). Above this mass region, the signal becomes marginal at SSC design luminosity. For this reason, the $H \rightarrow ZZ \rightarrow 2\ell + 2 \text{ jets}$ and $H \rightarrow WW \rightarrow \ell\nu + 2 \text{ jets}$ decay modes were also studied (their branching ratios are 20 and 150 times larger than that of the $H \rightarrow 4\ell$ mode). The signal to background ratio is much

Table 2-8

Examples of the expected physics performance of the SDC detector described in this Report. The results displayed here assume data samples corresponding to an integrated luminosity of 10 fb^{-1} , or one year at the SSC design luminosity of $10^{33} \text{ cm}^{-2}\text{s}^{-1}$. Any exceptions are explicitly noted.

Physics Process	Mass Region (GeV)	Physics Signature
Associated Higgs Production		
	80 - 150	$W + H, t\bar{t} + H \rightarrow \ell\gamma\gamma$
Direct Higgs Production		
	130 - 180	$H \rightarrow ZZ^* \rightarrow 4\ell$
	180 - 800	$H \rightarrow ZZ \rightarrow 4\ell$
	500 - 800	$H \rightarrow ZZ \rightarrow 2\ell 2\nu$
High Mass Boson Pairs		
Requires integrated luminosity of at least 50 fb^{-1} for complete studies	1-2 TeV	$Z\gamma \rightarrow \ell^+\ell^-\gamma$ $W^+Z \rightarrow \ell^+\ell^+\ell^-\nu$ $W^+W^+ \rightarrow \ell^+\ell^+$
Discovery of t Quark	$\lesssim 1 \text{ TeV}$	$t\bar{t} \rightarrow W^+W^- + X \rightarrow e^\pm\mu^\mp + X$
Mass Measurement of t Quark		
Sequential Dilepton Mode	$\lesssim 500$	$t\bar{t}$, one $t \rightarrow Wb$; $W \rightarrow e\nu$; $b \rightarrow \mu + X$ the other $t \rightarrow 3 \text{ Jets}$
Lepton + Jets + b -tag Mode	$\lesssim 500$	$t\bar{t}$, one $t \rightarrow W + X$; $W \rightarrow \ell\nu$ the other $t \rightarrow Wb \rightarrow b + 2 \text{ Jets}$
Non-standard t Decays		
Violation of τ Universality	$M_H \lesssim M_{\text{top}} - 15$	$t \rightarrow H^\pm b$; $H^\pm \rightarrow \tau^\pm\nu$; $\tau^\pm \rightarrow \pi^\pm + X$
Peak in 2-Jet Mass Distribution	$M_H \lesssim M_{\text{top}} - 25$	$t \rightarrow H^\pm b$; $H^\pm \rightarrow c\bar{s}$
Gluino and Squark Searches		
Missing- E_t + Jets	300 - 1000	$\tilde{g}\tilde{g} \rightarrow E_t^{\text{miss}} + 3-6 \text{ Jets}$
Like-Sign Dileptons	200 - 2000	$\tilde{g}\tilde{g} \rightarrow \ell^\pm\ell^\pm + 4 \text{ Jets}$
New Z Searches		
Discovery	$\lesssim 4 \text{ TeV}$	$Z' \rightarrow \ell^+\ell^-$
Width and Asymmetry	$\lesssim 2 \text{ TeV}$	$Z' \rightarrow \ell^+\ell^-$
Compositeness	$\Lambda \gtrsim 25 \text{ TeV}$	Inclusive Single Jet Spectrum

less favorable, due to the large contributions from the $W/Z + \text{jets}$ and $t\bar{t}$ processes. Nevertheless, these modes could provide an additional method for studying the very heavy Higgs region, allowing searches to be extended into the TeV region.

Following these studies in the context of the Minimal Standard Model, it is natural to explore what happens in more general models of the symmetry breaking sector. A more complex, but theoretically attractive, model is the minimal supersymmetric version of the Standard Model (MSSM). In this model, there are five Higgs bosons: three neutral (h^0, H^0, A^0), and two charged (H^\pm). The theory has two fundamental parameters and the analysis is more complex. It appears that over much of the parameter space, at least one of the neutral Higgs bosons should be visible, either in the SDC detector, or at LEP-II. However, some regions of the parameter space remain inaccessible.

The previous discussion focussed on the W^+W^- and ZZ final states, where the Higgs appears directly as a resonance. It is also important to study other boson pair channels to probe the electroweak theory more thoroughly. In particular, if no Standard Model Higgs is found below 1 TeV in mass, it is almost

certain that the symmetry breaking sector is strongly coupled. In this case, in analogy with QCD, one might expect resonant analogues of the ρ and ω to appear in the WW , WZ , and $Z\gamma$ channels (see Fig. 3-37). The discovery of such resonances is well within the capabilities of the SDC detector. The strong breaking could also manifest itself in non-resonant channels such as W^+W^+ , where the strong coupling would produce an excess of events over the Standard Model predictions. Further work remains to demonstrate convincingly that one can reduce the large like-sign $t\bar{t}$ background to the W^+W^+ final state to a manageable level. Preliminary studies, using lepton isolation and topological cuts, appear promising. The excellent lepton-charge measurement ability of the SDC detector plays a crucial role in removing the large opposite-sign backgrounds to the like-sign WW signal.

Physics of the t quark

The t quark is one of the few remaining ingredients of the Standard Model that has not been directly observed. Its existence is crucial to the Standard Model, and precision electroweak measurements restrict its mass, in the context of the Minimal Standard Model, to lie in the range $90 < M_{\text{top}} < 200$ GeV. The discovery and study of this quark thus allows stringent tests of this model.

The CDF and D0 experiments at Fermilab may very well discover the t quark in the next several years. If its mass is 150 GeV, these experiments may hope to reconstruct a handful of events. However, the SDC detector will be capable of reconstructing approximately 10^7 $t\bar{t}$ pairs in a year of SSC running. Such data samples will allow an accurate determination of the t quark mass by one of several methods. We have studied the mass measurement using the sequential decay of the t quark to an isolated electron (from the W decay) and a non-isolated muon (from the b -decay product of the same t quark). The dilepton mass spectrum provides a useful estimate of the t quark mass, with an estimated uncertainty, which is predominantly systematic, of 3 GeV after one SSC year (see Fig. 3-44). A second study has also been performed, using a lepton tag for one t quark decay and then reconstructing the mass of the three-jet system arising from the decay of the recoiling t quark via $t \rightarrow Wb \rightarrow b + 2 \text{ jets}$, where the b -jet is tagged in the SDC tracking system (see Fig. 3-51). This method would have a statistical error of 100 MeV after one year of running, but suffers from systematic uncertainties on the jet energy scale. The mass of the two non- b jets provides a clean W peak which can be used for a calibration, and we estimate the remaining systematic error on the mass measurement to be 3 GeV.

In addition to the mass measurement, it is important to study as many other properties of the t quark as possible. For example, in supersymmetric and other non-minimal models of the Higgs sector, the t quark can decay to a charged Higgs instead of a W . The charged Higgs will then decay to a tau lepton ($H^+ \rightarrow \tau\nu$) or to two jets ($H^+ \rightarrow c\bar{s}$). The branching fractions depend on the values of the parameters in the model. The former process manifests itself as a violation of lepton universality in t decays (see Fig. 3-56). The latter appears as a second peak in the mass distribution for the two non- b jets (see Fig. 3-66). In the context of the Minimal Supersymmetric Standard Model, if the charged Higgs is at least 20 GeV lighter than the t quark, it is possible for the SDC detector to observe its effects in t decays for any value of the remaining free parameter in the model.

SUSY searches

Supersymmetry has many theoretical attractions. For this discussion, we assume that SUSY particles must be produced in pairs, and that the lightest supersymmetric particle (LSP) is stable and neutral, and therefore behaves like a heavy neutrino. These assumptions are natural in many theoretical models. We have considered two basic discovery signatures that we can use to search for squarks and gluinos (the supersymmetric partners of ordinary quarks and gluons).

The first signature involves jets and missing transverse energy. The expected final state involves 3-6 jets plus missing transverse energy arising from the missing LSP. The cross section for gluino pair production is large, allowing the exploration of masses up to 1 TeV. The most difficult case, due to the small expected missing transverse energy, is for a relatively light gluino with mass ≤ 300 GeV (see Fig. 3-70). The SDC detector should be capable of finding a gluino in the mass range of 300 GeV to 1 TeV.

The second signature involves like-sign dileptons and jets. The like-sign dilepton signature is sensitive to a wider range of gluino masses than is the missing transverse energy signature. We have considered final states with two like-sign leptons and four additional jets, and estimate that a gluino in the mass range of 180 GeV to 2 TeV could be found in this manner (see Fig. 3-73). Furthermore, this signature provides a useful method to estimate the gluino mass with a precision of 10%.

Heavy boson searches

Models that enlarge the gauge group of the Standard Model predict the existence of additional gauge bosons. We have considered new Z bosons that arise in E_6 models (these models are theoretically very popular, and arise in many grand unification schemes, *e.g.*, “superstrings”). We have studied the properties of the new Z bosons that arise in such models, concentrating on their decays to lepton pairs, where we can measure the mass, width, cross section, and forward/backward asymmetry. Two extreme models within the E_6 family have been chosen, and new Z bosons with masses of 800 GeV and 4 TeV have been studied. The former provides a striking signal of more than 10^4 events per SSC year in the final state $Z \rightarrow \ell^+ \ell^-$ (see Fig. 3-75), while the later is at the limit of observability, giving some tens of events per SSC year (see Fig. 3-76). The detector resolution for the two-electron final state is adequate to extract the width of a new Z , and hence deduce some information about its couplings. Studies of the forward/backward asymmetry are possible in both the electron and muon pair final states, again providing strong separation between different models.

Compositeness

If quarks are made of more fundamental objects with a binding scale of order Λ , then one expects an enhancement of the inclusive jet cross section, relative to QCD predictions, at large values of the transverse momentum. In order to observe this effect, one fits to the shape of the inclusive jet spectrum for small transverse momentum and then extrapolates to large values to look for an excess. This technique places stringent requirements on the linearity of the jet energy measurement. With the jet linearity expected in the SDC calorimeter, and with proper single-particle calibrations, the systematic errors can be controlled, and the measurement is limited by statistics, leading to a bound on Λ of about 25 TeV after one SSC year.

QCD tests

When the SSC first turns on, we expect that our initial physics priorities will include Standard Model processes with large cross sections. We have surveyed the expected rates for single and multiple jet production, heavy quark cross sections and distributions, as well as single and multiple gauge boson production. The study of these theoretical predictions plays a particularly important role in understanding the expected backgrounds in the more exotic processes which the SDC detector will study.

References:

1. SDC Detector Parameters, SDT-000010, Rev. D, (April 1, 1992).

3. Physics and detector requirements

3.1. Overview

The role of this section is to survey the SSC physics relevant to the SDC detector, attempting to isolate the most demanding aspects of that physics, and derive the corresponding detector requirements. In doing so, it is appropriate to retain a certain scepticism towards the specific details of the physics analyses. One can reasonably expect some progress in our understanding during the intervening years, and our concept of exactly what physics is relevant may change somewhat. In striving for as much generality in the present analysis as possible, certain physics processes are used as archetypes to study, in a somewhat more abstract fashion, the capabilities that a general purpose SSC detector should possess.

In carrying out these studies, simple detector models have been used to try to capture the essentials of the behavior of the SDC detector in the SSC environment. More detailed considerations, and interactions with engineering designs, are confined to the subsequent detector sections. The design process for so complex a detector is highly iterative, moving back and forth between physics, instrumentation, engineering, and economics. The following work captures the collaboration at some intermediate point in that process, and attempts to render the process in a more logical and sequential fashion than is strictly correct. In that spirit, the models frequently used are parametrizations of the current baseline detector configuration in which some key parameters are varied in order to quantify the impact on the physics performance and capabilities.

In most plots in this section, there are far more Monte Carlo events than the number of events expected in a typical data run at the SSC. The error bars, however, correspond to the statistical errors on the expected numbers of events for the integrated luminosities defined in the figure labels, usually one “SSC year” of 10^{40} cm^{-2} or 10 fb^{-1} unless otherwise stated. Pileup events for the calorimeter simulations are included assuming an SSC luminosity of $10^{33} \text{ cm}^{-2}\text{s}^{-1}$ and a sensitive time of a single beam crossing. To the extent that the response time of individual detectors is slower than this, additional pileup events would need to be included.

3.1.1. Detector models

Tracking system

The performance of the tracking system is parametrized in terms of a resolution for charged tracks as a function of their 3-momenta. The parametrizations have been derived from the behavior of the baseline detector with the aid of a simulation of the measurement errors, misalignments, and multiple scattering in the material present in the tracking volume, as described in Chapter 4. Figure 3-1 displays the expected resolution as a function of pseudorapidity (η) for various values of p_t , where a beam constraint has been used to achieve the best possible resolution. We note that with the performance indicated here, the tracking resolution is better than the EM calorimeter resolution for p_t values less than about 100 GeV. The baseline tracking detector covers a pseudorapidity region of $|\eta| < 2.5$. The efficiencies for finding and identifying isolated individual tracks have been taken to be constant, independent of the details of the event topology.

Trigger system

The trigger system should be capable of selecting leptons and jets over the required angular range with an efficiency which is perfect above some threshold in transverse momentum. For lepton triggers, it is assumed that a single lepton trigger exists over the rapidity range $|\eta| < 2.5$ with a threshold of $p_t > 40$ GeV. A dilepton trigger with a threshold of $p_t > 20$ GeV for each lepton is also assumed for multilepton physics. Single and multiple jet triggers over the full calorimeter acceptance are also required for many generic studies, but are not directly addressed in this section. Finally, a missing transverse energy trigger with a threshold of 100 GeV is needed to select hard neutrinos or weakly interacting SUSY particles such as the photino. More complete evaluations of the expected performances of these triggers are contained in Chapter 8.




FIG. 3-1. The resolution of the baseline tracking system as a function of η for several p_t values. The solid curve is for $p_t = 1000$ GeV, and the dashed (dotted) is for $p_t = 250$ (100) GeV.

Calorimeter system

The calorimeter is assumed to cover the angular region $|\eta| < 6$ with a basic transverse segmentation that is a multiple of 0.05 in $\Delta\eta$ and $\Delta\phi$. The unit in phi is actually $\Delta\phi = 2\pi/128 \approx 0.049$. In the simulations described in the remainder of this section, the segmentation will be varied within the range of 0.05 to 0.8, but keeping a constant segmentation as a function of eta. The current electromagnetic (EM) calorimeter design has an η -dependent segmentation that increases by a factor of two from a base value of 0.05 at η values of 1.8, 2.6, 4.4, and 5.2. The current hadronic (HAD) calorimeter starts with a segmentation of 0.1 and increases by a factor of two at η values of 2.2, 4.4, and 5.2. The calorimeter is longitudinally segmented into an EM segment (EM) and two hadronic segments (HAD1 and HAD2), with depths given in Table 3-1. Parametrizations of the longitudinal and transverse distributions of energy deposited in individual calorimeter cells have been derived from EGS Monte Carlo simulations for electrons and ZEUS test beam data [1] for hadrons. These parametrizations are subsequently used, in conjunction with the single particle resolutions and nonlinearities, to simulate the response of the calorimeter in the following sections. A uniform magnetic field of 2.0 T is assumed to exist inside the barrel region.

In designing the SDC calorimeter, complex tradeoffs have been made between EM and hadronic single-particle response. High performance EM calorimetry demands fine sampling and large sampling fraction. Maintaining this sampling throughout the hadron calorimeter would be prohibitively expensive. The resulting discrepancy between the sampling in the EM and HAD calorimeters induces a π/e response ratio different from unity. There are other factors which further enhance this nonuniform response (choice of absorber material, ratio of absorber to scintillator thickness, etc.). A careful analysis of the physics requirements has been an essential ingredient in optimizing the SDC calorimeter, and the resulting design places greater emphasis on EM than on hadronic calorimetry.

The single particle resolution has been parametrized in terms of a stochastic term (a) and a constant term (b):

$$\frac{\sigma(E)}{E} = \frac{a}{\sqrt{E}} \oplus b.$$

The symbol \oplus means that the two terms are added in quadrature. This model can be generalized to the case where the sampling plates in the calorimeter are not projective, assuming that the stochastic term is due to the sampling, and hence varies as the square root of the effective plate thickness. The parametrization for the barrel ($|\eta| < 1.4$) is:

$$\frac{\sigma(E)}{E} = \frac{a}{\sqrt{E_t}} \oplus b$$

and for the endcap and forward ($1.4 < |\eta| < 6.0$) is:

$$\frac{\sigma(E)}{E} = \frac{a}{\sqrt{E_t}} \oplus b ,$$

where E_t is the transverse component of the energy (in GeV) and E_l is the longitudinal component of the energy (in GeV). All calorimeter models presented in the subsequent sections use this resolution model in which the barrel plates are parallel to the beam, and the endcap plates are perpendicular to the beam.

For the EM calorimeter designs currently under consideration for the SDC detector, plate thicknesses of 2-4 mm for the barrel and 2-8 mm for the endcap are being considered. To explore the performance implications, we have examined families of resolutions covering the range of 7.5% to 20% for the stochastic term and 0, 0.5%, 1.0%, and 2.0% for the constant term. The resolution for the baseline detector is given in Table 3-1. This corresponds to a design with 4mm lead plates in the barrel, 6mm lead plates in the endcap, and 4mm scintillator throughout. A higher performance option has also been considered. In this case, the stochastic term is taken to be 0.09 for the barrel and 0.14 for the endcap. This corresponds to the performance expected from 2mm lead plates in the barrel and 4mm lead plates in the endcap, keeping the 4mm scintillator throughout. The baseline constant term of 0.01 has been retained. We also assume that the electron response is linear, as detailed studies have shown that with a massless gap correction for energy deposited just behind the SDC coil, it is possible to achieve a linear response for energies above roughly 10 GeV (see Chapter 6).

The single particle response for hadrons has been parametrized from CALOR89 simulations. The stochastic and constant terms are displayed in Table 3-1. For the reasons outlined above, the hadron calorimeter is noncompensating with a resulting π/e response ratio as a function of energy that is parametrized as follows:

$$\pi/e = \alpha - \frac{\beta}{E^{0.15}} .$$

The coefficients in this expression have been derived from CALOR89 simulations. In order to fully define the coefficients, it is necessary to specify the calibration scheme. Here, we assume that the calorimeter has been calibrated in such a way that $\pi/e = 1$ at an incident energy of 300 GeV. This results in a π/e response that is less than one at lower energies and greater than one at higher energies. Preliminary test beam results indicate that the performance of the actual calorimeter may be slightly better than that given in Table 3-1 (smaller stochastic and constant terms in the resolution and a better π/e response ratio have been observed).

Muon system

The performance of the muon system is described in terms of a momentum resolution that is a function of p_t and η . The parametrization used here is shown in Fig. 3-2 for several values of muon p_t . This resolution has been derived from the covariance matrix for fits to simulated measurement points, ignoring any pattern recognition effects, but including the effects of multiple scattering and misalignments between different elements of the tracking system. The reconstruction efficiency is taken to be the same as that of the tracking system. The baseline design has a coverage in pseudorapidity of $\eta < 2.5$. More detailed studies are presented in Chapter 7.

Table 3-1

A summary of the parameters of the baseline SDC calorimeter which have been assumed in the subsequent analyses. The calorimeter depth is quoted in interaction lengths (λ).

Parameter	Barrel	Endcap	Forward
Coverage	$ \eta < 1.4$	$1.4 < \eta < 3.0$	$3.0 < \eta < 6.0$
Radius of front face (m)	2.10		
z position of front face (m)		4.47	12.00
Compartment depth			
EM (+ Coil)	1.1	0.9	
HAD1	4.1	5.1	13.0
HAD2	4.9	6.0	
EM resolution			
a	0.14	0.17	0.50
b	0.01	0.01	0.05
HAD resolution			
a	0.67	0.73	1.00
b	0.06	0.08	0.10
HAD nonlinearity			
α	1.13	1.16	1.16
β	0.31	0.38	0.38




FIG. 3-2. The resolution of the combined baseline tracking and muon system as a function of η for several p_t values. The solid curve is for $p_t = 1000$ GeV, and the dashed (dotted) curves are for $p_t = 250$ (100) GeV.

Lepton and photon identification

We take the global electron and muon efficiencies within the detector acceptance to be 85% for analyses requiring isolated leptons. This can be compared with CDF experience, where a value of $85 \pm 3\%$ is obtained for W and Z electrons, including the effects of triggering and mild isolation cuts[2]. In the case where the analysis requires two such leptons reconstructing to an on-shell Z boson, the lepton identification cuts are relaxed for the second lepton, and the efficiency for the second lepton is taken to be 95%. For electrons, this efficiency includes the effects of track finding and fitting as well as electron identification (e.g., an E/p

cut, Shower Max Detector profile cut, HAD/EM leakage cut, *etc.*). For muons, the efficiency includes track finding in both the inner tracker and the muon system, as well as a matching efficiency for associating the two detectors. Further details of the individual efficiencies are discussed in the relevant detector chapters.

When considering backgrounds to leptons or photons, it is necessary to assess the performance of the combination of all systems in the detector which contribute to rejecting the QCD jet backgrounds that are the major sources of fake electrons and photons. This is extremely difficult to simulate properly, and we have chosen to use measured values from current hadron collider experiments with similar identification strategies as a reference. For the identification of photons, the major source of background is the fragmentation of a jet into a leading neutral meson (π^0 , η^0 , or K_L^0). We have used the current rejection factor obtained by CDF in their single photon analysis[3], which is here defined as the ratio of background in their inclusive single photon cross section to the total inclusive jet cross section at the same p_t . This gives a rejection of roughly 5×10^{-4} against jets after a strict isolation cut of $E_t < 2$ GeV in a cone of $R < 0.7$ around the photon (where $R^2 = (\Delta\eta)^2 + (\Delta\phi)^2$). This is valid in the p_t region above about 20 GeV. For the identification of electrons, the major source of background is either the overlap between a charged pion track and a neutral meson in a jet, or a charged pion which interacts in the calorimeter in such a way that it looks very electromagnetic. In addition, conversion electrons arising from neutral meson decay or from photons interacting in the material in the tracking volume, although they are real electrons, are not primary electrons and hence must also be considered a background. We have again turned to CDF, using the ratio of the background in their inclusive electron sample (after conversion subtraction) to the total inclusive jet cross section at the same p_t . This gives a rejection factor of slightly better than 10^{-5} , or a factor of 50 better than that obtained for photon identification.

Jet definition

Jets are defined from the transverse energy deposition in each cell. Seed towers are found by summing over the three longitudinal calorimeter compartments using the coarsest segmentation of the three compartments (*e.g.*, the SDC baseline calorimeter with 0.05 EM and 0.1 HAD towers uses a seed tower with a size of 0.1). These seed towers are required to pass a threshold cut of 5 GeV in transverse energy, and then serve as the axis for a cone of the specified radius (typically $R = 0.6$ is used). All of the towers whose centers fall inside the cone are summed, using a threshold of 0.1 GeV in transverse energy. The momentum for the jet is then defined by vectorially adding all of the individual vectors pointing to the center of each tower in the jet. The jet energy is defined as the scalar sum of the individual tower energies. Using this definition, the jet acquires a mass, as though it were a massive particle decaying to a large number of massless particles. The original information about the directions of the individual particles in the jet is degraded by binning the energy depositions into calorimeter cells.

Using the ISAJET generator to create QCD two-jet events, and passing them through a detector with the properties described in Table 3-1, one obtains an energy resolution for jets in the barrel calorimeter of $0.61/\sqrt{E_t} \oplus 0.016$. The major feature of interest is the decrease in constant term, relative to that observed for single pions, due to the averaging which takes place when measuring the response to the many particles contained in the jet.

3.2. Electroweak symmetry breaking

In the simplest form of the Standard Model, the masses of the W and Z bosons, as well as the masses of the charged fermions, result from their interactions with a fundamental scalar field whose vacuum expectation value is not zero. One component of the scalar field manifests itself as the Higgs boson. The couplings of this boson are completely determined once the masses of the fermions and gauge bosons in the theory are known. The expected cross sections for producing the Higgs boson at the SSC via several different mechanisms are shown in Fig. 3-3.

Two mechanisms dominate the production of Higgs bosons in hadron-hadron collisions: the process $gg \rightarrow H$ which proceeds via a virtual quark loop[4], and the process $qq \rightarrow qqH$ via the exchange of two W or Z bosons[5]. Since the coupling of the Higgs to a quark is proportional to the quark mass, the production

rate from the first process is dependent on the t quark mass. Figure 3-3 shows the Higgs cross section that results from these processes at the SSC. It can be seen from this figure that the current allowed range of t -quark masses between the lower limit of 91 GeV from CDF [6] and the upper limit of 200 GeV inferred from precision tests of the standard model [7] produces an uncertainty of approximately a factor of four in the Higgs production cross section. In the remainder of this work we shall use a t -quark mass of 150 GeV.

If the mass of the Higgs boson exceeds 180 GeV, it is expected to decay predominantly to W^+W^- and ZZ pairs. At lower masses, the pattern of branching ratios is more complex, as indicated in Fig. 3-4. In making this figure we have taken into account the QCD corrections [8] which act to reduce the decay width for $H \rightarrow b\bar{b}$ and hence increase the branching ratio to other decay modes. Note that below the kinematic threshold for WW or ZZ production, the decay can still proceed through virtual W or Z production. This figure shows that the branching ratio to ZZ^* has a dip near $M_{\text{Higgs}} = 2M_W$ where the decay channel $H \rightarrow WW$ opens, and then falls rapidly as the Higgs mass is reduced below 140 GeV. The branching ratio to $\gamma\gamma$ is small but significant for Higgs masses below 160 GeV.




FIG. 3-3. The cross section for the production of a Higgs boson in pp collisions at $\sqrt{s} = 40$ TeV as a function of the Higgs boson mass for several different production mechanisms: gg fusion (solid), WW/ZZ fusion (dotted), $t\bar{t} + H$ production (dot-dashed), $W + H$ production (upper dashed), and $Z + H$ production (lower dashed). When the cross section depends on the t -quark mass, several curves have been included for different values of M_{top} .




FIG. 3-4. The branching ratio for a Higgs boson into various channels as a function of its mass.

In the following, we present a systematic analysis of Higgs discovery strategies and their impact on the SDC detector design. The first sections cover the Minimal Standard Model Higgs in three different mass regions (the strategies are different for each). Following that, there is a brief description of the impact of a more general Higgs sector (that of the Minimal Supersymmetric Standard Model) on the strategies developed for the Standard Higgs. Finally, an overview of the phenomenology of a strongly interacting Higgs sector is presented, with an example of how to search for it in the W^+W^+ channel.

3.2.1. Low-mass Higgs ($80 < M_{\text{Higgs}} < 130$ GeV)

The LEP-II exclusion region for the Standard Model Higgs is expected to extend up to roughly 80 GeV [9], using the $e^+e^- \rightarrow H + Z \rightarrow b\bar{b}\nu\bar{\nu}$ process. This process results in about 10 detected events for 500 pb^{-1} of data (several years of running at the $10^{31} \text{ cm}^{-2}\text{s}^{-1}$ design luminosity), and is expected to be the most powerful channel for Higgs discovery. Note also that the cross section for $e^+e^- \rightarrow ZZ$ is several times larger than the signal process for a Higgs mass of 80 GeV, so that it is necessary to use kinematic constraints to improve the $b\bar{b}$ mass resolution in order to separate a possible signal from the nearby background. These complications make it unlikely that LEP-II will be able to improve its mass reach, and hence hadron colliders will need to extend their sensitive range down into the $80 < M_{\text{Higgs}} < 140$ GeV range in order to conclusively explore the Standard Model.

Previous discussions of this region concentrated on direct Higgs production followed by the rare decay $H \rightarrow \gamma\gamma$. This decay mode has the advantage that it does not suffer from the profuse Z background that will exist at the SSC, and the branching ratio $\text{BR}(H \rightarrow \gamma\gamma)$ is relatively constant over the Higgs mass range of interest. Unfortunately, the QCD backgrounds to the direct production mechanism, $q\bar{q} \rightarrow \gamma\gamma$ and $gg \rightarrow \gamma\gamma$, are very large, requiring a detector with extraordinary photon energy resolution [10] in order to see the signal. In addition, there are vast numbers of π^0 's from QCD 2-jet production which require a very high quality preshower detector in order to achieve the necessary rejection factor of at least 10^7 against QCD 2-jet backgrounds.

In the SDC LoI, an alternative solution was presented, relying on the associated production of a W boson and the Higgs [12], followed by the decay $H \rightarrow \gamma\gamma$. In this case, one uses the leptonic decay $W \rightarrow \ell\nu$ to tag the event, and then requires two additional photons arising from the Higgs decay. An extensive background analysis was performed, indicating that with an accumulated luminosity of 10^{41} cm^{-2} , a significant signal would be visible above the backgrounds which arise predominantly from non-resonant $W + \gamma\gamma$ and $W + \gamma + \text{jet}$ processes. Meanwhile, the process $gg \rightarrow t\bar{t} + H$ was re-discovered, long after the original calculations had been performed [13]. For the heavier t -quark masses now current, this process provides a substantial increase in the expected "high p_t isolated lepton" + Higgs rate (a factor of roughly 5 at the SSC, somewhat less at the LHC) over the $W + H$ process alone. Several phenomenological analyses have already appeared in the literature [14], demonstrating that the expected improvement is relatively insensitive to the t -quark mass, provided it is in the range $100 < M_{\text{top}} < 200$ GeV. In the remainder of this work we shall use a t -quark mass of 150 GeV. We also assume that the t quark decays with the expected Standard Model branching ratio to leptons.

In the present section, we first review the signal which would be expected in the SDC baseline detector for direct Higgs production followed by the decay $H \rightarrow \gamma\gamma$. The marginal nature of this signal leads naturally into the analysis of the more complex associated production processes $W + H$ and $t\bar{t} + H$. For associated production, we first summarize the current situation for calculations of the signal and background processes. Next, full event simulations using PYTHIA for the two signal production mechanisms $W + H$ and $t\bar{t} + H$ are described. The intent is to demonstrate that these events still allow identification and isolation of the leptons and photons with reasonable efficiency, as well as reconstruction of a narrow peak in the two-photon invariant mass distribution, despite the complex nature of the final states. Finally, parton level calculations of the signals and backgrounds are used to explore the acceptances and to look for kinematic cuts which can suppress backgrounds. This leads into a discussion of the expected significance of the signal and implications for detector performance requirements.

Direct Higgs production

The direct production of a Higgs boson in the intermediate mass region followed by its decay to $\gamma\gamma$ has a relatively large cross section at the SSC. For the best case, a Higgs mass of 150 GeV, there are several hundred events after experimental cuts. Unfortunately, these events sit on top of a large irreducible background from the QCD production of photon pairs via the processes $q\bar{q} \rightarrow \gamma\gamma$ and $gg \rightarrow \gamma\gamma$. This is evident in Fig. 3-5, which indicates the continuum of events expected from these background sources in addition to the narrow peaks expected from Standard Model Higgs signals with masses as defined in the caption. In addition, there is a very large background arising from QCD 2-jet production, where the two jets are mis-identified as photons. Reducing this background source to a level below that of the two photon background requires a rejection factor for both jets of at least 10^7 , which is somewhat beyond the SDC detector performance discussed in Section 3.1.1. Although it is clearly important, we have neglected this background in the subsequent discussions. Furthermore, the precise mass reconstruction required for this final state implies knowledge of the event vertex at the millimeter level. For events with multiple interactions within a beam crossing, there is no unambiguous way of associating charged tracks which contain precise vertex information with the particular interaction which produced the Higgs.




FIG. 3-5. The two-photon invariant mass distribution including the signals from Higgs bosons with masses of 80, 100, 120, 140, and 160 GeV. The background includes only the irreducible backgrounds arising from the $q\bar{q} \rightarrow \gamma\gamma$ and $gg \rightarrow \gamma\gamma$ processes.

In order to explore the ability of the SDC detector to observe this final state, studies have been performed including only the real $\gamma\gamma$ backgrounds mentioned above. Two different calorimeter models have been used, one with the SDC detector baseline performance given in Table 3-1, and the second being the high performance option described in Section 3.1.1. This latter option has a stochastic term of 9% in the barrel and 14% in the endcap, combined with the baseline constant term of 1%. The two photons were required to satisfy acceptance cuts of $p_t > 20$ GeV and $|\eta| < 2.5$. To further suppress the backgrounds, a cut on the angle of the photons in the two-photon center of mass, $|\cos\theta^*| < 0.8$, was also made. The precision of the photon angular measurements was assumed to be 1 mrad, corresponding to the expected precision of the SDC Shower Maximum detector of 2 mm (and assuming that the event vertex can be located unambiguously). We note that this particular analysis was performed at the particle level only—no account was taken of actual shower shapes or realistic energy reconstruction algorithms in the calorimetry.

The resulting signals are displayed in Fig. 3-6, where the irreducible backgrounds shown in Fig. 3-5 have been subtracted using the statistical errors expected after one year of data at SSC design luminosity. This optimistic estimate of the signal assumes that no systematic errors are present in the background subtraction. Signals have been included at Higgs masses of 80, 100, 120, 140, and 160 GeV. The signal remaining after the background subtraction is rather marginal, even with the optimistic background

treatment assumed here. For the baseline case, the statistical significance, evaluated by fits to the signal plus background, is roughly five sigma for a mass of 140 GeV. This rises to seven sigma for the high performance case shown in Fig. 3-7. These numbers suggest the possibility of confirming a Higgs in this channel for masses of roughly 120 to 150 GeV, but further analysis of the jet backgrounds is required before this claim could be made with confidence.




FIG. 3-6. The expected signals from Higgs bosons with masses of 80, 100, 120, 140, and 160 GeV. The irreducible backgrounds displayed in Fig. 3-5 have been statistically subtracted using an exponential fit to the background shape. The baseline calorimeter performance has been assumed.




FIG. 3-7. The same as Fig. 3-6 except that the high performance calorimeter option described in the text has been assumed.

Associated Higgs production

In studying the associated production of the Higgs boson, we have used the matrix elements for the $W + H$ and $t\bar{t} + H$ processes to study the signal. For the background analysis, we have included contributions from those irreducible backgrounds which already contain the two photons as part of the matrix element. In addition, estimates have been made of the contributions from events where at least one of the photons in the matrix element is replaced by a jet. Such contributions are then reduced by the experimental rejection factor of 5×10^{-4} given in Section 3.1.1. Although the required matrix elements for heavy quark production ($Q\bar{Q} + \gamma + jet$ and $Q\bar{Q} + 2 jets$) are not yet available, the dominant source of jets in these final states is from the decays of the $Q\bar{Q}$ system itself. We have therefore used the matrix elements for $Q\bar{Q} + \gamma$ and used the jets produced in the Q decays as candidates for the second photon. Similarly, for the case where both photons arise from mis-identified jets, inclusive $W + jets$ and $t\bar{t}$ events are the dominant background sources.

Physics simulation

The results in this section were produced using the PYTHIA[17] Monte-Carlo event generator (version 5.5) and the partonic generator PAPAGENO (version 3.60). The PYTHIA generator contains only the signal processes, and uses the original calculation of Kunszt[13] to implement the $q\bar{q} \rightarrow t\bar{t} + H$ and $gg \rightarrow t\bar{t} + H$ production mechanisms. Its advantage is in using phenomenological strings in the parton-shower framework to give a detailed description of all produced particles in a hadron collider event.

The PAPAGENO generator includes all of the processes relevant to the present study. It uses an independent calculation by Ian Hinchliffe for the $t\bar{t} + H$ matrix elements. The $W + \gamma\gamma$ background process uses a matrix element calculation from Mangano[12]. The $Q\bar{Q}\gamma\gamma$ backgrounds are implemented using the

calculation of Ballestrero and Maina[15]. The $QQ\gamma$ backgrounds are implemented using the calculation of Ellis and Kunszt[16]. This generator contains a reasonably accurate model for the sequential decay of heavy quarks to leptons, but does not provide an accurate estimate for the flow of hadronic energy arising from the quark fragments or the underlying event and so cannot be used to estimate the effects of isolation or lepton identification requirements. It also does not incorporate any branching ratios for Higgs decay, so we have chosen to use the Kunszt and Stirling results[8], including the $\mathcal{O}(\alpha_s^2)$ QCD corrections which result in a running b -quark mass. These corrections significantly reduce the rate for Higgs decay to b quark pairs, resulting in an increased branching ratio to two photons in the intermediate mass region. The PYTHIA calculation also includes these QCD corrections in an approximate form.

The $t\bar{t} + H$ cross section increases slowly with M_{top} ($\pm 15\%$ variation for $100 < M_{\text{top}} < 200$ GeV). In contrast, the $t\bar{t}\gamma\gamma$ background is strongly dependent on M_{top} (varying in the ratio 9:3:1 for $M_{\text{top}} = 100, 150, 200$ GeV). We have adopted the nominal value of $M_{\text{top}} = 150$ GeV. As this background gives the dominant contribution, should M_{top} really be as low as 100 GeV, the overall background would be increased by roughly a factor of three. It is also worth noting that the background processes include only those contributions obtained by attaching photon lines to the primary partons (W , b , or t). Potentially important contributions can also arise where one photon is attached to a primary parton and the other is produced in a radiative decay such as $W \rightarrow \ell\nu\gamma$ or $t \rightarrow Wb\gamma$. These contributions will be most significant for low photon p_t (small compared to the mass of the decaying object). Detailed calculations are required to estimate their contribution in the kinematic region studied in the present analysis.

For the $W + H$ signal process, the higher order QCD corrections are known (they are essentially the same as those for direct W production), and the result is a small (10% to 15%) increase at SSC energies[18]. For the $t\bar{t} + H$ process, the QCD corrections are unknown, but the dependence of the result on the q^2 scale is known to be large. The common (and conservative) choice used in the literature is $q^2 = \hat{s}$, where \hat{s} is the $t\bar{t}H$ center of mass energy. This choice gives a mean value for q of roughly 700 GeV. The scale used here is $q^2 = \langle M_t^2 \rangle$, where $M_t^2 = p_t^2 + M^2$ is computed for each of the three primary particles ($t\bar{t}H$) in the matrix element, and then the average of these three values is used. This scale gives a mean q value of roughly 200 GeV, and a cross section which is 30% larger than that obtained with $q^2 = \hat{s}$. We have chosen the smaller scale because for other processes (e.g., QCD 2-jet production, which has now been calculated to $\mathcal{O}(\alpha_s^3)$), scales proportional to p_t^2 are known to be more reasonable than scales proportional to \hat{s} .

Kinematics and acceptance

The kinematics of the signal events have been studied in order to understand how best to optimize detection. The two signal processes, particularly for light Higgs masses, are quite different. For the $W + H$ process, the mass of the W is quite small on SSC scales, so that the required \hat{s} is also quite small. This leads to a very broad rapidity distribution and a rather narrow p_t distribution. For the $t\bar{t} + H$ process, the necessity of producing two heavy quarks in addition to the Higgs implies that its rapidity distribution is much narrower. The Higgs rapidity distribution is reflected in the pseudo-rapidity distribution of the lepton and photons, leading to very different acceptances for the two processes. For the $W + H$ process, good acceptance requires coverage out to $\eta = 5$, whereas the $t\bar{t} + H$ process is well suited to the $\eta = 2.5$ coverage of the SDC detector.

The transverse momentum distributions of the lepton and photons are only slightly different for the two processes. From Fig. 3-8 for the $t\bar{t} + H$ process with a Higgs mass of 80 GeV, it appears that a requirement of $p_t > 20$ GeV is acceptable for the photons, and it is possible to use a higher lepton- p_t requirement without significant losses. For the $W + H$ process with a Higgs mass of 80 GeV, the softer p_t spectrum makes the acceptance smaller by a factor of two for the p_t requirement of 20 GeV on the lepton and photons. For a larger Higgs mass of 160 GeV, both processes have good acceptances for the $p_t > 20$ GeV requirement. We note that for the $t\bar{t}H$ process, decreasing the pseudorapidity coverage from 2.5 to 1.5 decreases the signal acceptance by a factor of 1.8, whereas increasing the coverage from 2.5 to 3.0 increases the acceptance by about 10%.




FIG. 3-8. Families of acceptance curves for $t\bar{t} + H$, with $M_{\text{Higgs}} = 80$ GeV (dashed) and $M_{\text{Higgs}} = 160$ GeV (solid).

- (a) The fraction of events with at least one lepton with $p_t > p_0$ as a function of p_0 . The lepton has $|\eta| < 1.5$ (lowest curve), 2.0 (lower middle curve), 2.5 (upper middle curve), or 3.0 (upper curve).
 (b) The fraction of events with at least two photons with $p_t > p_0$ as a function of p_0 . The photons have $|\eta| < 1.5$ (lowest curve), 2.0 (lower middle curve), 2.5 (upper middle curve), or 3.0 (upper curve).

The lepton acceptance for $t\bar{t}$ events has been checked by comparing PYTHIA and PAPAGENO simulations of this final state. PYTHIA is observed to give a much softer p_t spectrum for the lepton from the W decay. For a t -quark mass of 150 GeV, the mean p_t from PYTHIA is roughly 55 GeV, whereas that from PAPAGENO is roughly 75 GeV. This is because of the different physics models used. PAPAGENO decays the t quark directly into a W and a b quark (ignoring any QCD radiation that might accompany the decay), giving a p_t distribution which may be slightly too hard. PYTHIA instead fragments the t jet into normal mesons and baryons plus a t meson or baryon, which subsequently decays to a W and a b , etc. This complex model is not expected to work well for large t -quark masses, and will give a significantly softer lepton p_t spectrum. The result is that 38% of the PAPAGENO $t\bar{t}$ events have a W decay lepton with $p_t > 20$ GeV, whereas only 32% of the PYTHIA events have such a lepton. For a higher cut of $p_t > 40$ GeV, PAPAGENO gives 24% and PYTHIA gives 22%. The PAPAGENO result has been used in the present analysis, as it is expected to be more accurate for $M_{\text{top}} = 150$ GeV. We further note that for $p_t > 20$ GeV, 70% of the highest- p_t leptons in $t\bar{t}$ decays arise from W decays, whereas for $p_t > 40$ GeV, this number rises to 85%.

Lepton and photon identification and measurement

The associated production of the Higgs with other massive particles is expected to produce significantly more complex events than the clean $H \rightarrow ZZ \rightarrow 4\ell$ events which are frequently studied as a performance reference. This section focuses in particular on the issues related to identifying and measuring electrons and photons in the calorimeter, as this is where the effects of this additional complexity may be felt most strongly. Reconstructing non-isolated muons is much easier because of the lack of confusion outside the calorimeter.

In considering the issue of electron and photon identification and measurement, it is useful to identify two different types of requirements. The first is the very local (*i.e.*, roughly the size of the EM shower itself) region where the actual identification and energy measurement take place. Any excess energy in this "identification" region will both degrade the identification capability and influence the energy measurement. One of the most critical distributions for electron and photon identification is the transverse shower profile measured by the Shower Maximum detector. Here, due to the fact that the Shower Maximum detector represents a thin slice through the electromagnetic shower, the intrinsic fluctuations are large. Estimates from CDF data indicate that additional energy deposits of less than about 1 GeV will have a negligible

effect on the shower profile for a 50 GeV electron or photon (one will most likely use Shower Maximum information from 7-9 strips to define the profile, which corresponds to an area in $\Delta\eta \times \Delta\phi$ of roughly 0.05 by 0.05). The more demanding requirement is likely to be the energy measurement, where a 50 GeV electron or photon is expected to be measured with 1-2% precision, and hence the pileup effects need to be kept to less than 1% of the electron or photon p_t . A measurement of the energy to this level of precision in the SDC detector requires an area in $\Delta\eta \times \Delta\phi$ of roughly 0.1 by 0.1, and therefore is more sensitive to pileup than the transverse shower profile. In the following discussions, we will concentrate on the energy measurement. A further complication arises at larger values of the pseudo-rapidity where the calorimeter segmentation in pseudorapidity is decreased in order to maintain a reasonable physical cell size (the cell size is roughly $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ at $\eta = 2.0$ and 0.2×0.2 near $\eta = 3.0$). In this case, the electron or photon measurements will require a larger region in (η, ϕ) space, and hence the pileup will be larger.

Second, once a lepton (electron or muon) or photon is identified, it is often necessary to further reduce backgrounds by requiring that the amount of energy observed near the lepton or photon be small. This "topological" isolation helps eliminate leptonic backgrounds arising from decays of heavy quarks and photon backgrounds from π^0 's in QCD jets. Studies described in Section 3.2.2 have shown that a cut of 5 GeV of excess transverse energy in a cone of radius $R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} = 0.3$ around the lepton or photon is very effective at reducing backgrounds from leptons produced in b or c quark decays from $t\bar{t}$ events, providing a rejection of 10 to 1000, depending on the p_t of the lepton. In the present analysis, the background from the $b\bar{b}\gamma\gamma$ process would be unacceptably large without the lepton isolation requirement. Monte Carlo studies for photons indicate similarly large rejections can be achieved against photons associated with jets.

In order to proceed with a more detailed analysis, it is necessary to define an energy measurement algorithm. This allows a precise study of the effects of both energy reconstruction and isolation in associated Higgs production. The minimum region required to measure the energy of an EM shower is clearly a single EM tower, but this will badly underestimate the true energy when the shower approaches the edge of a cell due to energy leakage into adjacent cells. Here, we consider two different algorithms — one using a 2×2 array of cells, the other using a 3×3 array. In order to collect the energy which leaks from the EM calorimeter into the HAD1 calorimeter, we also include all those cells which lie behind the EM towers used in the array (if the center of a HAD1 cell lies behind an EM cell in the array, it is used). The direction of the electron is defined from the track impact point in the calorimeter, whereas that of the photon is defined using the Shower Maximum detector. The calorimeter impact point defined in this manner is used to choose the best 2×2 or 3×3 array of cells. The energy of the electron or photon is then defined by summing the EM and HAD1 energies in the cell array.

Once the electron or photon energy has been defined, the isolation variable is defined to be the excess E_t deposit in the cone of radius R (E_t in the cone minus the reconstructed E_t of the electron/gamma), rather than a variable which is scaled by the particle p_t . The argument is that the excess energy in the cone for signal events should be roughly constant (ignoring remnants of the EM shower subtraction), whereas that for the background should increase with particle p_t . Therefore, the fixed cut should have a constant efficiency and a monotonically improving rejection as a function of the p_t . The distribution of excess E_t expected for signal events is shown in Fig. 3-9. It is apparent that the $t\bar{t} + H$ process does produce significantly fewer isolated photons than does the cleaner $W + H$ process. The signal efficiency for several different isolation cone sizes is summarized in Table 3-2. The global efficiency of these rather strict isolation cuts is far from unity, but their background rejection is very high.

Table 3-2

A summary of the observed efficiencies of energy isolation cuts for events arising from associated Higgs production versus the size of the isolation cone. These samples used $M_{\text{Higgs}} = 80$ GeV, but studies with $M_{\text{Higgs}} = 160$ GeV give identical results. An excess transverse energy of less than 10 GeV was required (this cut could be optimized for each cone size). The photons and leptons were required to have $|\eta| < 2.5$ and $p_t > 20$ GeV.

Particle	Radius	$W + H$ process	$t\bar{t} + H$ process
Highest p_t ℓ from W	0.2	0.99 ± 0.003	0.94 ± 0.008
	0.3	0.97 ± 0.004	0.87 ± 0.01
	0.4	0.95 ± 0.006	0.80 ± 0.01
Highest p_t ℓ from b	0.2		0.13 ± 0.02
	0.3		0.058 ± 0.013
	0.4		0.035 ± 0.010
Either γ from Higgs	0.2	0.99 ± 0.003	0.95 ± 0.007
	0.3	0.97 ± 0.004	0.90 ± 0.009
	0.4	0.96 ± 0.005	0.82 ± 0.01
Isolation on all lepton/photons	0.2	0.97 ± 0.004	0.86 ± 0.01
	0.3	0.93 ± 0.006	0.73 ± 0.01
	0.4	0.89 ± 0.01	0.58 ± 0.02




FIG. 3-9. The distribution of excess E_t in a cone of radius $R = 0.3$ around the photon direction for photons from associated Higgs production. The photons were required to have $|\eta| < 2.5$ and $p_t > 20$ GeV. (a) for photons from the $W + H$ process with $M_{\text{Higgs}} = 80$ GeV. (b) for photons from the $t\bar{t} + H$ process with $M_{\text{Higgs}} = 80$ GeV.

Mass resolution

The energy resolution of the EM calorimeter is the major factor influencing the invariant mass resolution for the two photons from the Higgs boson decay. A series of studies have been done at the particle level, simply smearing the energy of each photon by the expected calorimeter resolution. The mass resolution as a function of Higgs mass for several relevant cases is summarized in Table 3-3. Several features are apparent. First, for the stochastic terms of the baseline design, the constant term also plays a significant role. It is

desirable to keep it below the 1% level. Second, in order to take full advantage of the improved stochastic terms characteristic of the high performance option, it is desirable to have a constant term of 0.5%.

Table 3-3

Two photon invariant mass resolutions in GeV for events from the $t\bar{t} + H$ process. The entries in the table are the sigma of a Gaussian fit to the signal (in GeV). The simulation was done at the particle level using parametrized resolutions, where "Base" refers to the terms given in Table 3-1. The final columns summarize the resolution expected for the high performance option defined in Section 3.1.1.

M_{Higgs}	$a = \text{Base}$ $b = 0\%$	$a = \text{Base}$ $b = \text{Base}$	$a = \text{Base}$ $b = 2\%$	$10\%/\sqrt{E}$	$a = 9\%/14\%$ $b = 0.5\%$	$a = 9\%/14\%$ $b = \text{Base}$
80	1.08	1.23	1.56	0.67	0.80	0.93
100	1.24	1.44	1.89	0.78	0.93	1.11
120	1.39	1.65	2.19	0.87	1.05	1.28
140	1.52	1.81	2.51	0.96	1.16	1.44
160	1.64	2.00	2.81	1.03	1.25	1.61

Using the photon energy measurement algorithm discussed previously, it is possible to study in detail the effects of segmentation and pileup on the mass resolution. A series of simulations have been performed for a fixed Higgs mass of 80 GeV, using segmentations between 0.05 and 0.2, and luminosities between $10^{33} \text{ cm}^{-2}\text{s}^{-1}$ and $5 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$. Several features are apparent from this study; in Fig. 3-10, we attempt to summarize them. First, a region of 0.1×0.1 in (η, ϕ) is not sufficient to contain the energy of an EM shower for $|\eta| > 2.0$ (the SDC calorimeter design changes the EM cell size from 0.05 to 0.1 at $\eta = 1.8$). Second, the effects of shower spreading and pileup in the energy measurement region induce an additional smearing in the mass resolution of roughly 0.4 GeV, thereby increasing the parametrized result of 1.1 GeV for the SDC baseline calorimeter to about 1.2 GeV. Third, the resolution is not a strong function of HAD1 segmentation, even at elevated luminosities of $5 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ (the mean value of the true minus observed mass changes by less than 250 MeV and the resolution deteriorates by less than 5% for a HAD1 segmentation of 0.2). Finally, the resolution is a strong function of EM segmentation, and $\Delta\eta \times \Delta\phi$ of 0.1 is the coarsest segmentation which gives acceptable performance (the mean value of the true minus observed mass shifts by 800 MeV and the resolution deteriorates by 20% at five times design luminosity). Figure 3-10 indicates that even for this segmentation, the mass resolution is developing a non-Gaussian tail at higher luminosity. These studies support the SDC baseline choice of 0.05 EM segmentation, and indicate that the performance of such a calorimeter is adequate for precision mass measurements well beyond the SSC design luminosity.

Significance of signal

Finally, the signals from the two production mechanisms are combined with the events from the different background processes to give the mass plots shown in Fig. 3-11 and Fig. 3-12. The $b\bar{b}\gamma\gamma$ background has been reduced by a factor ten to account for the isolation cut applied to the lepton. The plots contain the expected data sample obtained after one year of running at the SSC design luminosity. The actual contributions to the signal peaks are itemized in Table 3-4 and Table 3-5. An additional requirement that all leptons and photons be separated from each other by at least $\Delta R = 0.4$ has been added. This cut is very useful for reducing $b\bar{b}\gamma\gamma$ backgrounds with negligible loss of signal events.

From these plots, it is possible to extract the statistical significance of the expected result. The proper way to do so involves performing a polynomial fit to the background, excluding the signal region. These background parameters, plus a Gaussian for the signal, are then fit to the full distribution. An analysis of the integral of the Likelihood function leads to the best estimate of the significance of the signal. Such an




FIG. 3-10. The observed two photon invariant mass resolution as a function of the calorimeter segmentation and luminosity. The variable plotted is the true $\gamma\gamma$ mass minus the reconstructed $\gamma\gamma$ mass. (a) Mass resolution for a 2×2 cell algorithm. The solid curve is for 0.05 EM and HAD1 segmentation at design luminosity. The dashed curve is for 0.05 EM segmentation and 0.2 HAD1 segmentation at five times design luminosity. The dot-dashed (dotted) curve is 0.1 (0.2) EM and HAD1 segmentation at five times design luminosity. (b) Mass resolution for a 3×3 cell algorithm.

Table 3-4

A summary of the number of events produced in and detected by the SDC detector per SSC year as a function of the Higgs mass. The detected events are required to satisfy $p_t > 20$ GeV and $|\eta| < 2.5$. A $\Delta R = 0.4$ cut between all of the leptons and photons was also applied. A factor $(0.85)^3 = 0.614$ has been used to account for lepton and photon efficiency. An additional factor of 0.93 for $W + H$ events and 0.73 for $t\bar{t}H$ events has been applied to account for the isolation requirement of less than 10 GeV of excess E_t in a cone of radius $R = 0.3$ (see Table 3-2). All results are derived from a simulation using PAPAGENO, $\text{BR}(H \rightarrow \gamma\gamma)$ as given in Ref. 8, and $\text{BR}(t \rightarrow \ell + X) = \text{BR}(W \rightarrow \ell + X) = 0.22$ (only W leptons are accepted from the t quark decay—all others will fail the isolation cuts).

M_{Higgs}	$W + H$ events Produced	$W + H$ events Detected	$t\bar{t} + H$ events Produced	$t\bar{t} + H$ events Detected
80	25.0	4.0	49.6	12.2
100	22.0	4.2	48.9	13.7
120	18.4	3.7	45.5	13.7
140	10.2	2.2	28.0	8.8
160	1.6	0.4	5.0	1.6

analysis is lengthy, and must be performed on many statistically independent samples in order to guarantee the detection of the (eagerly awaited) signal. Here, a standard naive approximation is made, namely the background is estimated, and the significance is defined to be $d = \text{signal}/\sqrt{\text{background}}$. This is a measure of the probability that a Gaussian distributed background sample could have fluctuated up to impersonate a signal. If one wants to include some estimate for the background uncertainty, the background value used in the calculation can be chosen to be the largest reasonable estimate.

The significance d defined in this manner is a useful criterion for evaluating the relative performance of different detector options, although a value greater than five is likely to be necessary for discovery. This

Table 3-5

A summary of the number of signal and background events observed within two sigma of the peak for different Higgs masses, where sigma is the expected mass resolution for the given Higgs mass. The event sample is identical to that described in Table 3-4. The baseline sample is shown in Fig. 3-11 and the high performance sample is shown in Fig. 3-12.

M_{Higgs}	Signal events Baseline resolution	Background events Baseline resolution	Signal events High performance	Background events High performance
80	15.6	7.7	15.8	6.2
100	17.2	7.5	17.5	6.2
120	16.9	7.3	17.2	6.3
140	10.7	6.6	10.6	4.9
160	1.9	5.6	1.9	4.3




FIG. 3-11. The two-photon invariant mass distribution for the expected signals from Higgs bosons of mass 80, 100, 120, 140, and 160 GeV. The background curves are cumulative, and are (from lowest to highest): $b\bar{b}\gamma\gamma$, $t\bar{t}\gamma\gamma$, and $W\gamma\gamma$. The baseline detector resolution has been assumed.




FIG. 3-12. The two-photon invariant mass distribution for the expected signals from Higgs bosons of mass 80, 100, 120, 140, and 160 GeV. The background curves are cumulative, and are (from lowest to highest): $b\bar{b}\gamma\gamma$, $t\bar{t}\gamma\gamma$, and $W\gamma\gamma$. The high performance calorimeter option defined in Section 3.1.1 has been assumed.

definition has very simple scaling properties, namely the significance decreases as the square root of the resolution for fixed integrated luminosity, and increases as the square root of the integrated luminosity for fixed resolution. This means that a factor of two improvement in resolution will do the same physics in one half the elapsed time. Figure 3-13 and Fig. 3-14 display the results of this simple calculation for several families of resolution functions. From these figures one can conclude that with the baseline stochastic term (see Table 3-1), a constant term of 1% is an appropriate goal. If the stochastic term can be reduced to 9%, then a reduction in the constant term to 0.5% is desirable (see also Table 3-3).




FIG. 3-13. A series of curves showing the expected significance of the combined $W + H$ and $t\bar{t} + H$ signal above the background for one SSC year of integrated luminosity for different Higgs masses. Events from both $W + H$ and $t\bar{t} + H$ have been included. For the solid curves, the stochastic term is assumed to be that of the baseline detector (see Table 3-1), and the constant term is taken to be, from highest to lowest, 0.5, 1.0, or 2.0%. For the dotted curves, the stochastic term is assumed to be that of the high performance option defined in Section 3.1.1, and the constant term is taken to be 0.5 and 1.0%. Finally, an additional dashed curve with $10\%/\sqrt{E}$ resolution is included for reference.




FIG. 3-14. A series of curves showing the expected significance of the combined $W + H$ and $t\bar{t} + H$ signal above the background for one SSC year of integrated luminosity for different Higgs masses. Events from both $W + H$ and $t\bar{t} + H$ have been included. The constant term is fixed to be 1% and the four curves represent, from highest to lowest, stochastic terms of 7.5%, 10%, 15%, and 20%.

3.2.2. Intermediate-mass Higgs ($130 < M_{\text{Higgs}} < 180$ GeV)

As long as the Higgs boson mass is below the threshold for W^+W^- and ZZ decays, it remains very narrow (Γ_{Higgs} is less than 100 MeV for $M_{\text{Higgs}} < 160$ GeV). The most spectacular signature in the intermediate mass region is the decay $H \rightarrow ZZ^* \rightarrow 4\ell$ (here $*$ indicates a virtual particle). The branching ratios for the Standard Model Higgs shown in Fig. 3-4 indicate that this mode has a reasonable branching ratio for Higgs masses above 130 GeV. This provides good overlap with the $\gamma\gamma$ mode which is useful below 150 GeV.

Physics simulation

Studies of the signal for direct Higgs production followed by the decay $H \rightarrow ZZ^*$ have been performed using the WW/ZZ fusion and gg fusion processes contained in the PYTHIA generator. Note that in all calculations made by this version of PYTHIA, the Z 's are allowed to be off-shell as demanded by the available phase space, so that ZZ^* or even Z^*Z^* production will automatically be generated. The ZZ continuum background is also generated using PYTHIA. This process includes the contribution from both the Z and the γ propagators, and hence also contains the process $Z\gamma^*$ in the approximation that the γ^* is radiated from a quark line (this is relevant for small M_{Z^*}). The PAPAGENO generator was used to efficiently generate the leptons arising from heavy flavor decays ($t\bar{t}$, $Z + b\bar{b}$, and $Z + t\bar{t}$) which provide additional multi-lepton backgrounds.

Acceptance and kinematics

Events are selected by requiring two leptons with transverse momenta greater than 20 GeV. This results in a trigger efficiency of $\sim 50\%$ (70%) for a Higgs mass of 120 (160) GeV. Figure 3-15 shows the acceptance of the detector as a function of the minimum lepton transverse momentum for several values of the rapidity coverage (the default coverage is $|\eta| < 2.5$). In this study we assume that leptons of p_t less than 10 GeV cannot be detected with adequate background rejection. Hence it can be seen that requiring that the four leptons be in the SDC detector fiducial volume and pass the triggering (2 leptons with $p_t > 20$ GeV) and detection (4 leptons with $p_t > 10$ GeV) requirements results in an efficiency of 17% (42%) for a Higgs mass of 120 (160) GeV. This acceptance would decrease by roughly a factor of 2.5 if the lepton coverage were restricted to $|\eta| < 1.5$, and would increase by 25% if the coverage were increased to $|\eta| < 3.0$.

To simplify the analysis, only the four highest- p_t leptons inside the acceptance are used in the subsequent calculations (this also provides some further rejection against backgrounds which produce more than four leptons). The charge-zero lepton pair with mass closest to M_Z is required to satisfy $M_{\ell\ell} = M_Z \pm 10$ GeV. The other lepton pair is required to have $M_{\ell\ell} > 20$ GeV to reduce the $q\bar{q} \rightarrow Z\gamma^*$ background.

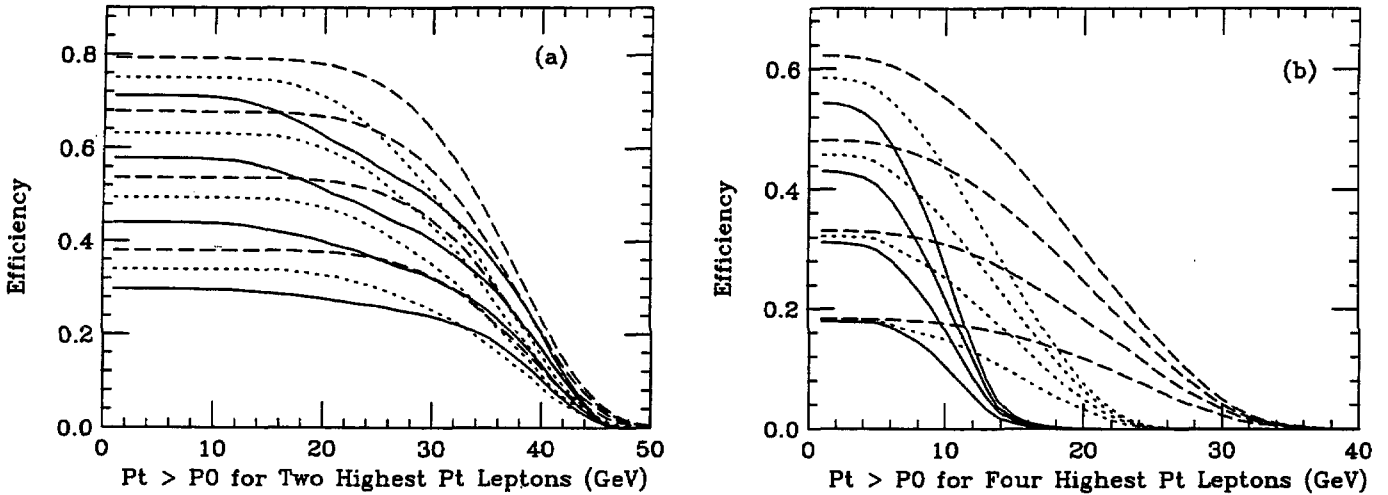


FIG. 3-15. Families of acceptance curves for $H \rightarrow ZZ^*$, for $M_{\text{Higgs}} = 120$ GeV (solid), 140 GeV (dotted), and 160 GeV (dashed). (a) The fraction of events with at least two leptons with $p_t > p_0$ as a function of p_0 . Both leptons have $|\eta| < 1.5$ (lower curve), 2.0 (lower middle curve), 2.5 (upper middle curve), or 3.0 (upper curve). (b) The fraction of events containing two leptons with $p_t > 20$ GeV and $|\eta| < 2.5$ plus two others with $p_t > p_0$ and $|\eta| < 1.5$ (lower curve), 2.0 (lower middle curve), 2.5 (upper middle curve), or 3.0 (upper curve).

Lepton identification

In order to study the background rejections and signal efficiencies as a function of the lepton isolation requirement, we have used a sample of $t\bar{t}$ and $H \rightarrow 4\ell$ events from PYTHIA. The isolation requirement has been studied for cone sizes varying from $R = 0.2$ to $R = 0.4$. The definition of this quantity has been given in the Section 3.2.1. The results of this study, summarized in Table 3-6, show the signal efficiency and background rejection for different bins in lepton p_t . The lepton class was defined based on the process and the identity of the lepton's parent. In this table, only electrons have been included since the additional EM shower energy they produce inside the isolation cone makes them more difficult to isolate than muons.

Further insight into these efficiencies can be gained by looking at the distribution of E_t in the isolation cone for the different classes of leptons. Figure 3-16 shows distributions of the isolation variable for the lowest p_t bin defined in this study ($10 < p_t < 20$ GeV) for a cone of $R = 0.3$. The first distribution in the set of four is for leptons coming from a sample of $H \rightarrow ZZ^*$ events, assuming $M_{\text{Higgs}} = 140$ GeV. The other three distributions are for leptons coming from a large sample of $t\bar{t}$ events, where the parent was a

Table 3-6

A summary of the observed efficiencies for detecting electrons with different parents arising from two different processes as a function of the electron p_t and the radius of the surrounding cone. Electrons are accepted if there is less than 5 GeV of excess transverse energy in this cone. The “Z parent” column refers to electrons arising from $H \rightarrow ZZ^*$ with $M_{\text{Higgs}} = 140$ GeV. The other columns are for $t\bar{t}$ events with $M_{\text{top}} = 150$ GeV. Some entries are 68% confidence limits based on no observed events.

p_t range	Radius	Z parent	W parent	b parent	c, d, u parent
$10 < p_t < 20$	0.2	0.98 ± 0.005	0.96 ± 0.01	0.27 ± 0.01	0.034 ± 0.004
	0.3	0.95 ± 0.007	0.88 ± 0.01	0.11 ± 0.008	0.010 ± 0.002
	0.4	0.90 ± 0.01	0.79 ± 0.01	0.046 ± 0.005	0.005 ± 0.002
$20 < p_t < 30$	0.2	0.98 ± 0.003	0.94 ± 0.01	0.13 ± 0.01	0.015 ± 0.004
	0.3	0.96 ± 0.01	0.88 ± 0.01	0.045 ± 0.007	0.004 ± 0.002
	0.4	0.89 ± 0.01	0.79 ± 0.01	0.015 ± 0.004	0.001 ± 0.001
$30 < p_t < 50$	0.2	0.98 ± 0.003	0.94 ± 0.01	0.047 ± 0.008	0.013 ± 0.004
	0.3	0.93 ± 0.005	0.89 ± 0.01	0.0012 ± 0.004	0.006 ± 0.003
	0.4	0.87 ± 0.01	0.79 ± 0.01	0.005 ± 0.003	0.002 ± 0.002
$50 < p_t < 150$	0.2	0.96 ± 0.004	0.95 ± 0.005	0.006 ± 0.003	0.001 ± 0.001
	0.3	0.89 ± 0.01	0.87 ± 0.01	0.002 ± 0.002	< 0.001
	0.4	0.81 ± 0.01	0.77 ± 0.01	< 0.002	< 0.001

W, a b quark, or a c or lighter quark. This figure and Table 3-6 indicate a problem that occurs when trying to reject low p_t b quark backgrounds, namely the b decay products are not tightly collimated, and even a moderate cone size may often not include them, thereby leading to many “isolated” background leptons. This emphasizes the need for relatively large isolation cones for low- p_t leptons, placing stringent constraints on the allowed calorimeter noise contributions. As the p_t of the lepton increases, the p_t of the heavy parent also increases, so that the decay products are more energetic and more collimated, making the isolation requirement increasingly powerful. For p_t of the lepton in the lowest bin (mean p_t of order 15 GeV), rejections of roughly ten are possible using a cone of $R = 0.3$ while maintaining a high efficiency for true isolated leptons. When the p_t of the lepton has risen to 100 GeV, these rejections increase to almost a factor of 1000 (the present study lacked the statistical power to evaluate this factor in detail).

Mass resolution

For the SDC baseline detector, in the kinematic region relevant for $H \rightarrow ZZ^*$ decays, the parametrized tracking resolution is significantly better than the calorimeter resolution. For $M_{\text{Higgs}} = 140$ GeV, the four-muon mass resolution is almost identical to that of the four-electron mass measured with the tracking system, and is 0.8 GeV. The four-electron mass resolution measured with the calorimeter is 1.9 GeV for the same Higgs mass. For the muon final states, it is possible to take full advantage of this resolution. However, for final states with electrons, the parametric resolution is degraded by the passage of the electrons through the material in the tracking volume and the subsequent induced bremsstrahlung. A detailed simulation, described in Chapter 4, and summarized in Fig. 3-17, indicates the severity of the effect. For the calorimeter, even at a p_t of 10 GeV, the bremsstrahlung photon and the original electron will almost always lie within 10 cm of each other, and hence be included in the 2×2 cell array used to reconstruct the energy (although the track will no longer point exactly to the centroid of the energy cluster in the calorimeter). Thus, the calorimeter energy measurement is largely immune to the bremsstrahlung effects which degrade the tracking performance. It is possible that more sophisticated track reconstruction algorithms can be developed to cope with the bremsstrahlung problem, but in the present analysis we have




FIG. 3-16. The distribution of excess E_t in a cone of radius $R = 0.3$ for different classes of electrons. The electrons are all in the range $10 < p_t < 20$ GeV. (a) Electrons from $H \rightarrow ZZ^*$ for $M_{\text{Higgs}} = 140$ GeV. (b) W electrons coming from t quark decays ($M_{\text{top}} = 150$ GeV). (c) b electrons coming from t quark decays. (d) c (or u, d) electrons coming from t -quark decays.




FIG. 3-17. The invariant mass reconstructed by the SDC tracking system for ZZ^* decaying to four leptons for a Higgs mass of 140 GeV. The plots contain events which have been passed through a detailed simulation of the tracking system, and contain the effects of material in the tracking volume, which severely degrades the performance for electrons due to the induced bremsstrahlung. (a) the mass distribution for the $4e$ final state. (b) the mass distribution for the 4μ final state.

chosen to use calorimetric reconstruction for the electron pairs.

For the calorimetric reconstruction, a study of mass resolution using the energy reconstruction algorithms described in Section 3.2.1 has been performed. The results of this study, in which various segmentations were examined, are very similar to those for the $H \rightarrow \gamma\gamma$ case (the $H \rightarrow 4e$ case is more sensitive to pileup and segmentation since there are four measurements rather than two). The effect of the transverse shower size and pileup at design luminosity is to add a smearing of roughly 0.5 GeV in quadrature with the parametrized calorimetric resolution of 1.9 GeV for a Higgs mass of 140 GeV. For EM segmentation of 0.05 and HAD1 segmentation of 0.2, the mean value of the true mass minus the

reconstructed mass is shifted by 500 MeV and the mass resolution is degraded by 15% at five times design luminosity. If both EM and HAD1 segmentation are 0.1, the mean value of the true mass minus the reconstructed mass is changed by 1.8 GeV and the mass resolution is degraded by 25% at five times design luminosity.

Signal significance

For a Higgs mass of 140 GeV, there are approximately 250 produced events in final states with e 's or μ 's in one standard SSC year.

There are four potential sources of lepton backgrounds. Note that by requiring four identified leptons, all non-leptonic backgrounds become negligible. First, the production of ZZ^* or $Z\gamma^*$ final states from either $q\bar{q}$ or gg initial states can give rise to the 4ℓ final state. The rate from the $q\bar{q}$ process is small[19]. The rate from gg initial states is not known; we have increased the rate from $q\bar{q}$ by a factor of 1.65 to estimate its effect. The rate for the process $gg \rightarrow Z\gamma$ (with an on-shell photon) is known[20]; it is approximately 20% of that for $q\bar{q} \rightarrow Z\gamma$, and hence our background may be an overestimate. Our larger estimate is based upon comparisons of the calculated rates for $q\bar{q} \rightarrow ZZ$ and $gg \rightarrow ZZ$ at larger ZZ invariant masses (see Ref. 21). Next, there are backgrounds from final states containing a Z and a pair of heavy quarks ($Z + b\bar{b}$ or $Z + t\bar{t}$), where semileptonic decays of the t and b quarks give rise to additional leptons. While the cross section for the former process is much greater than for the latter, the leptons from t decay tend to be isolated while those from b decay are not. Finally, there is the contribution from the $t\bar{t}$ final state. Four semileptonic decays resulting from this final state will give rise to a potential background. Two of the leptons (from the decays of t and \bar{t}) are likely to be isolated, while the others are not. Hence the isolation requirement is very effective in removing this background.

The signal events have one pair of leptons whose invariant mass reconstructs to the Z mass. Unlike the other backgrounds, the $t\bar{t}$ final state does not have such a pair of leptons. Hence, it can be reduced by requiring that there be one pair of leptons whose invariant mass reconstructs to the Z mass. The background rejection achieved by this requirement will depend on the resolution of the detector. Once the resolution is comparable to the natural width of the Z no further improvements in resolution will reduce this background. A further rejection can be obtained by removing like-sign lepton pairs which are produced by the heavy flavor decays discussed above. This provides a reduction of roughly 1.5 for the $Z + b\bar{b}$ and $Z + t\bar{t}$ backgrounds, and roughly 2 for the $t\bar{t}$ background.

The isolation rejection defined in the previous section, requiring that the excess E_t in a cone of radius $R = 0.3$ be less than 5 GeV, has been applied to these backgrounds. We have used the conservative assumption that this requirement reduces the non-isolated backgrounds by a factor of ten, independent of the p_t of the lepton (see Table 3-6). This means that $t\bar{t}$ and $Z + b\bar{b}$ cross sections are reduced by a factor 10^2 (assuming that only 2 of the 4 leptons come from b or lighter quark decays), and the $Z + t\bar{t}$ cross section is left unchanged (assuming that all 4 leptons come from Z 's or W 's).

Applying the isolation requirement discussed above to all four leptons in the signal results in an additional loss of roughly 6% of the events for each isolated lepton (a total efficiency of 78% for all four leptons). Finally there is an efficiency for the identification and reconstruction of electrons and muons in the detector. This is assumed to be 85% for each lepton unless the leptons reconstruct to a Z , in which case we assume an efficiency of 95% for the second of the two leptons. The total effect of acceptance and efficiency is to reduce the number of events from 250 produced to roughly 40 observed for a Higgs of mass 140 GeV.

Figure 3-18 shows the reconstructed invariant mass distribution for Higgs bosons with masses in the range 120 to 170 GeV, together with the expected backgrounds, using calorimeter measurements for the electron momenta. Although the mass resolution for the four muon channel is twice as good as for the four electron channel, the signals in the different channels have been combined.

The statistical significance of the peaks shown in Fig. 3-18 has been evaluated by counting the number of events predicted above the expected backgrounds in the vicinity of each peak (a region of $\pm 2\sigma$ was




FIG. 3-18. The reconstructed Higgs mass for ZZ^* decaying to $4e$, 4μ , and $2e2\mu$ with $M_{\text{Higgs}} = 120, 130, 140, 150, 160$, and 170 GeV, including the expected backgrounds. The background curves are cumulative, and are (from lowest to highest): $q\bar{q} \rightarrow ZZ^*$, multiplied by 1.65 to account for $gg \rightarrow ZZ^*$, $Z + b\bar{b}$, $Z + t\bar{t}$, and $t\bar{t}$. The invariant mass has been calculated using calorimeter measurements for the electrons.

used, where the σ is an average resolution from a fit to the 4μ , $2e2\mu$, and $4e$ channels combined). The $M_{\text{Higgs}} = 130$ GeV through $M_{\text{Higgs}} = 170$ GeV peaks are unambiguous. There are 18 signal events over an expected 3.5 event background for 130 GeV, and 32 signal events over an expected 10 event background for 170 GeV for one year of SSC running at nominal luminosity. The $M_{\text{Higgs}} = 120$ GeV peak has one fifth the signal and one half the expected background when compared to the 130 GeV peak, and consequently requires several years of SSC running to establish a convincing signal. At lower masses the rates are very small due to the rapidly falling branching ratio into the ZZ^* final state (see Fig. 3-4). Our requirement that leptons have $p_t > 10$ GeV also results in an acceptance that falls as the Higgs mass is reduced. An increase in the integrated luminosity in excess of 10^{40}cm^{-2} will enable the mass range to be extended slightly downwards, but no significant result below about 120 GeV can be expected from this mode.

3.2.3. Heavy Higgs ($180 < M_{\text{Higgs}} < 800$ GeV)

A Standard Model Higgs whose mass is above the WW and ZZ thresholds decays predominantly into these final states. As the mass approaches 800 GeV, the width of the resonance increases dramatically: $\Gamma_{\text{Higgs}} = 1.4, 30, 269$ GeV for $M_{\text{Higgs}} = 200, 400, 800$ GeV. Furthermore, the small production cross section makes discovery more difficult. In this section, we first review the leptonic decay modes $ZZ \rightarrow \ell^+\ell^-\ell^+\ell^-$ and $ZZ \rightarrow \ell^+\ell^-\nu\bar{\nu}$, where the signals are relatively clean but statistically limited. In order to study the higher mass region with higher statistics, it is very attractive to try to use the hadronic decays of the W and Z bosons. The branching ratio for $H \rightarrow ZZ \rightarrow \ell^+\ell^- + 2 \text{ jets}$ is roughly 20 times that for $H \rightarrow ZZ \rightarrow 4\ell$, and the branching ratio for $H \rightarrow WW \rightarrow \ell\nu + 2 \text{ jets}$ is roughly 150 times that for $H \rightarrow ZZ \rightarrow 4\ell$. Due to the additional $W/Z + \text{jets}$ and $t\bar{t}$ backgrounds for these final states, the signal to background ratio is much less favorable than that for the purely leptonic decays. Nevertheless, there are two techniques that play an important role in rejecting the large backgrounds. The first involves reconstructing the mass and decay angular distribution for the $W/Z \rightarrow 2\text{-jet}$ decay, the second involves searching for the very forward jets produced by the WW/ZZ fusion process (Fig. 3-3 indicates that for large Higgs mass, this process dominates Higgs production at the SSC). These two techniques and their detector implications are discussed in detail.

$H \rightarrow ZZ \rightarrow 4\ell$

The gg fusion and WW/ZZ fusion processes have been studied using the PYTHIA generator. It should be noted that the s -pole approximation is used for the Higgs matrix elements (although an energy-dependent Breit-Wigner is used for the line shape), which is known to overestimate the very high mass tail of a heavy Higgs. It is claimed that PYTHIA overestimates this rate by about 10% for $M_{ZZ} < 1200$ GeV with an 800 GeV Higgs mass. A detailed discussion of a complete matrix element calculation, with comparisons to the approximations used here, may be found in Ref. 22.

In this section, we consider the specific cases of $M_{\text{Higgs}} = 200$ GeV, 400 GeV, and 800 GeV. Once we are above the ZZ threshold, the signal consists of two pairs of leptons each of which has an invariant mass consistent with the Z mass (defined as $M_Z \pm 10$ GeV). Again, we select events with two leptons of $p_t > 20$ GeV and require that all four leptons be isolated and have $p_t > 10$ GeV. The background arises from the processes $q\bar{q} \rightarrow ZZ$ [23] and $gg \rightarrow ZZ$ [21]. Backgrounds from the final states $Z + t\bar{t}$, $Z + b\bar{b}$, or $t\bar{t}$ have also been considered. It is apparent that these backgrounds are small once we require that there are two lepton pairs whose invariant mass reconstructs to $M_Z \pm 10$ GeV (the backgrounds are much lower than those in the intermediate mass case because we can require two lepton pairs close to the Z mass). They become negligible once we make the additional requirement that the four leptons be isolated.

When the Higgs mass is as large as 800 GeV, the Z bosons produced in its decay have a substantial Lorentz boost, and their decay products will frequently lie close together. Figure 3-19 shows the distribution in the minimum ΔR (the distance in (η, ϕ) space) between any two leptons in such events. This distribution has significant implications for the segmentation of the EM calorimeter. In order to identify and reconstruct the electrons in such cases, one will probably need at least one quiet cell separating the 2×2 or 3×3 cell arrays used for the reconstruction. This requires that the EM cells be smaller than 0.1 units in $\Delta\eta$ and $\Delta\phi$ in order to maintain high efficiency.




FIG. 3-19. The distribution of the minimum ΔR (the distance in (η, ϕ) space) between any two leptons in Higgs events with a Higgs mass of 800 GeV.

Figure 3-20 shows the reconstructed Z mass for lepton pairs produced in the decay of an 800 GeV Higgs. The charge-zero pair with invariant mass closest to the Z mass is selected as the first Z candidate, and the other lepton pair is taken to be the second Z candidate. Figure 3-20 contains both pair masses, plotted separately for the individual lepton types. Note that as the p_t of the leptons increases beyond roughly 100 GeV, the SDC tracking system has worse resolution than the SDC EM calorimeter. The figure indicates that the resolution for the SDC detector, as assumed in the baseline design, is a good match to the 10 GeV window around the Z mass used in this analysis. In more quantitative terms, the acceptance for the $M_Z \pm 10$ GeV requirement on both pairs for the 4μ channel (reconstructed using the combined

tracking and muon system) is reduced by 12% relative to that for the $4e$ channel (reconstructed using the EM calorimeter) due to this difference in resolution. The acceptance of the detector as a function of the minimum lepton transverse momentum for several values of rapidity coverage is shown in Fig. 3-21. The acceptance for the cuts used in the present analysis varies from 53% for a 200 GeV Higgs mass to 70% for an 800 GeV Higgs mass. It is evident that the p_t threshold has little influence on the acceptance, whereas increasing the η coverage from 2.5 to 3.0 would increase the acceptance by $\sim 20\%$. Decreasing the lepton coverage from $|\eta| < 2.5$ to $|\eta| < 1.5$ would decrease the acceptance by a factor of roughly 2.5.




FIG. 3-20. The invariant mass distribution for lepton pairs from $H \rightarrow ZZ$ events for a Higgs mass of 800 GeV. The tracking measurements are used for the muons, whereas the calorimeter measurements are used for the electrons.

(a) the distribution for the $4e$ final state. (b) the distribution for the 4μ final state.

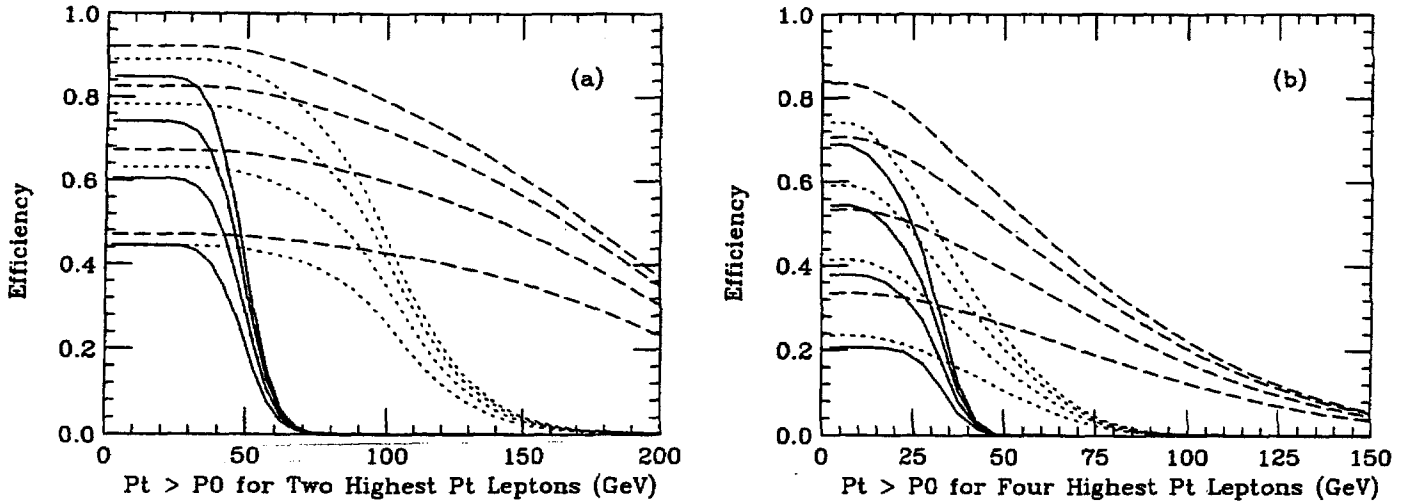


FIG. 3-21. Families of acceptance curves for $H \rightarrow ZZ$, for $M_{\text{Higgs}} = 200$ GeV (solid), 400 GeV (dotted), and 800 GeV (dashed).

(a) The fraction of events with at least two leptons with $p_t > p_0$ as a function of p_0 . Both leptons have $|\eta| < 1.5$ (lower curve), 2.0 (lower middle curve), 2.5 (upper middle curve), or 3.0 (upper curve). (b) The fraction of events containing two leptons with $p_t > 20$ GeV and $|\eta| < 2.5$ plus two others with $p_t > p_0$ and $|\eta| < 1.5$ (lower curve), 2.0 (lower middle curve), 2.5 (upper middle curve), or 3.0 (upper curve).

Figure 3-22 through Fig. 3-24 show the result of adding the signal and expected background for the three Higgs masses considered here. The electrons have been reconstructed from the calorimeter information

(giving better resolutions in this kinematic regime), whereas the muons use the combined tracking and muon system resolutions. The different resolutions in these channels do not manifest themselves in significantly different ZZ mass spectra, due to the large intrinsic width of the Higgs boson at these masses. The continuum background from $q\bar{q} \rightarrow ZZ$ has been multiplied by 1.65 to take account of the $gg \rightarrow ZZ$ process which is not included in the Monte-Carlos. In the case of Higgs masses of 200 and 400 GeV, a clear peak above background can be seen and the rates are sufficient for a discovery to be made. In the case of a Higgs mass of 800 GeV, there are fewer signal events and no clear peak. We note that the line shape for an 800 GeV Higgs is sufficiently broad (the intrinsic width is 270 GeV) that the mass distribution no longer peaks at 800 GeV, due to the falling parton flux. The peak of a relativistic Breit-Wigner with an energy dependent width will be shifted downwards by about 50 GeV by the exponentially falling parton luminosity.




FIG. 3-22. The ZZ invariant mass distribution showing a peak due to a Higgs of mass 200 GeV. The two lepton pairs were both required to have $M_{\ell\ell} = M_Z \pm 10$ GeV. The background curves have the same significance as those of Fig. 3-18, but the ZZ continuum background gives the only visible contribution.




FIG. 3-23. Same as Fig. 3-22, except that the Higgs mass is 400 GeV.




FIG. 3-24. Same as Fig. 3-22, except that the Higgs mass is 800 GeV.




FIG. 3-25. Same as Fig. 3-24, except that both Z 's were required to satisfy $p_t(Z) > 200$ GeV.

The signal to background ratio for the 800 GeV Higgs can be improved by requiring that the transverse momenta of the Z 's be more than 200 GeV. As is evident in Fig. 3-25, the background has been reduced with little loss in signal. The peak region contains 20 events with 6 expected background. To claim a signal in the 800 GeV case, one must have confidence in one's knowledge of the ZZ continuum rate at large ZZ invariant masses that would be expected in the absence of a Higgs boson. The measured ZZ rate at lower invariant masses can be used to reduce the uncertainties in the theoretical predictions for the ZZ rate at large values of the invariant mass. The major uncertainties in this extrapolation arise from the structure functions and higher order QCD corrections[24]. We estimate that we can determine the background with an uncertainty of 20%, and therefore the $M_{\text{Higgs}} = 800$ GeV signal would require 2-3 years of SSC running at the nominal luminosity to be sufficient for discovery.

At very large values of the Higgs mass, we are limited by event rate. In the case of the four lepton channel, we are confident that the effect of pile-up will not significantly modify the backgrounds displayed in Fig. 3-25 as the luminosity approaches $10^{34} \text{ cm}^{-2}\text{s}^{-1}$. Preliminary studies of the muon resolution, assuming that the silicon and outer layers of the central tracker as well as the full muon system function normally, indicate that the resolution in the relevant p_t range will deteriorate by only a small amount. This change in resolution can be easily compensated by slightly enlarging the window used for Z selection, with only small increases in the background.

$H \rightarrow ZZ \rightarrow \ell^+ \ell^- \nu \bar{\nu}$

In view of the limited rate for $M_{\text{Higgs}} = 800$ GeV, we have looked at the channel with e^+e^- or $\mu^+\mu^-$ and missing- E_t in order to extract the signal from the decay $H \rightarrow ZZ \rightarrow \ell^+ \ell^- \nu \bar{\nu}$. This decay is analogous to the $W \rightarrow \ell \nu$ decay, in which the transverse mass distribution shows a narrower peak than the Jacobian seen in the transverse momentum distribution (this occurs because the transverse mass also uses the information contained in the angle between the two decay products in the transverse plane). Here, the appropriate variable would be:

$$M_t^2 = \left\{ [p_t(Z)^2 + M_Z^2]^{1/2} + [(E_t^{\text{miss}})^2 + M_Z^2]^{1/2} \right\}^2 - p_t(Z)^2 - (E_t^{\text{miss}})^2 - 2p_t(Z)E_t^{\text{miss}} \cos \theta ,$$

where θ is the angle between the $Z \rightarrow \ell\ell$ vector and the missing- E_t vector in the transverse plane. The missing- E_t vector approximates the momentum carried by the Z which decayed to $Z \rightarrow \nu\bar{\nu}$. Because the Higgs peak has become extremely wide for $M_{\text{Higgs}} \geq 700$ GeV, the missing- E_t and M_t distributions look almost identical. We have chosen to use the missing- E_t , due partially to the simpler behavior of the backgrounds. The impact of the large Higgs width is shown in Fig. 3-26, where the signal in the missing- E_t distribution for $H \rightarrow ZZ \rightarrow \ell^+ \ell^- \nu \bar{\nu}$ is quite apparent, albeit less prominent than the signal in the invariant mass distribution for the $H \rightarrow 4\ell$ final state. The branching ratio for the former is six times larger than that for the $H \rightarrow 4\ell$ mode.

Just as for the four-lepton final state, one of the backgrounds is from continuum ZZ production. However, additional backgrounds arise from: $Z + \text{jets}$ final states where the jets are mis-measured or lost in the forward direction, from $Z + \text{heavy quarks}$ where the missing- E_t arises from semileptonic decays, and from $t\bar{t}$ events where two isolated leptons are produced that have an invariant mass close to the Z mass. The missing- E_t capability of our detector now plays a vital role in our ability to extract a signal. This is emphasized in Fig. 3-27, which displays the overwhelming background to the missing- E_t signature in a detector without a forward calorimeter. Additional discussions of the missing- E_t signature appear in Section 3.4.

The background from the $t\bar{t}$ final state is computed by requiring that there be a pair of isolated leptons (these arise almost exclusively from the decay $t \rightarrow b\ell\nu$) of opposite charge whose invariant mass is $M_Z \pm 20$ GeV. The resolution of the SDC detector is somewhat better than this but the limited Monte-Carlo statistics prevent us from making a tighter cut. In this case the missing- E_t is mainly due to neutrinos and not to resolution effects.




FIG. 3-26. The invariant mass distributions for $H \rightarrow 4\ell$ and twice the missing- E_t for the process $H \rightarrow 2\ell 2\nu$ are compared. The signal for the latter has been reduced by a factor of six so that the area of the signals is approximately the same. The intention of the figure is simply to show the shape of the signal distribution. The dot-dashed curve is for the 4ℓ final state; the dashed curve is for the $2\ell 2\nu$ final state. A Higgs mass of 800 GeV has been assumed.




FIG. 3-27. The comparison of the missing- E_t for the signal $H^0 \rightarrow ZZ \rightarrow \ell^+\ell^-\nu\bar{\nu}$ and the background from $Z + jets$ events in which one or more jets at $|\eta| > 3$ is lost out the end of the detector. This summarizes the effect of leaving out the forward calorimeter in SDC, which leads to substantial increases in backgrounds to processes with missing- E_t . A Higgs mass of 800 GeV has been assumed.

In the case of the $Z + jets$ background, there are two detector effects that contribute to the background. The jet resolutions assumed for our detector [25] have very long non-Gaussian tails (modeled on those observed in the CDF detector) and hence there is a non-negligible possibility that the jet will have its energy mis-measured by a very large amount and will give rise to a large apparent missing- E_t . These tails are due to cracks in the CDF detector as well as to neutrinos arising from semileptonic decays (mainly bottom and charm quarks). In order to reduce the background from the final state $Z + jets$ we have required that there be no jets in the event that are measured to have transverse energy above 300 GeV. This has a negligible effect on the signal but reduces this background by at least a factor of five for a missing- E_t of 300 GeV or greater. If the non-Gaussian tails from our jet resolutions are much less severe, we can drop this jet veto since the background from $Z + jets$ becomes insignificant provided the calorimetry coverage is adequate to prevent high p_t jets from being lost. The second detector effect is that caused by jets in the forward region being lost out of the end of the forward calorimeter. In this context, it is worth remarking that a jet with transverse momentum of 300 GeV at $\eta = 5$ has an energy of 22 TeV and hence is kinematically forbidden. The SDC baseline design includes fiducial coverage for $\eta < 5$ in order to eliminate this background source.

There are several possible final states involving a Z and heavy quarks. We have looked at the background from the $Z + t\bar{t}$ and $Z + b\bar{b}$ final states. There are two ways in which these processes can contribute. First, events can have $Z \rightarrow \nu\bar{\nu}$ decays and two leptons from semileptonic heavy quark decays, in which the two leptons have an invariant mass close to the Z (we require $M_Z \pm 20$ GeV). Second, events can have $Z \rightarrow \ell\ell$ and missing- E_t arising from the heavy quark decays. Summing these two contributions produces fewer than 3 events (without the jet veto) in the bin of Fig. 3-28 with missing- E_t of 300 GeV where there are about 30 signal events. The background contributions from $Z \rightarrow \ell\ell$ are the dominant ones, with the process $Z + t\bar{t}$ giving about 2 events and the process $Z + b\bar{b}$ giving one event. By including the jet veto, the t quark contribution does not change much but the bottom contribution falls by at least a factor of ten (this estimate is limited by Monte-Carlo statistics).

Figure 3-28 shows the missing- E_t distribution for events where we have required that there be an e^+e^-

or $\mu^+\mu^-$ pair with mass $M_Z \pm 20$ GeV and $p_t > 250$ GeV, and that there be no detected jets with $p_t > 300$ GeV. The two dominant detector dependent backgrounds are shown separately. It can be seen that these are negligible compared with the irreducible background from continuum production of the ZZ final state. Again, this method of searching for the Higgs is therefore limited by the production rate and not the background. The strategy for extracting a signal from this final state is similar to that in the 4-charged lepton channel. Measurements of that channel and of the missing- E_t spectrum at lower values are used to reduce the uncertainties in the predicted overall missing- E_t spectrum in the absence of a Higgs boson. The observed excess of events, in conjunction with results from the 4ℓ channel, could provide strong evidence for the existence of the Higgs boson within one SSC year.

In the case of the $H \rightarrow \ell^+\ell^-\nu\bar{\nu}$ final state, experiments at higher luminosity are likely to be affected by the need to modify or replace the calorimetry for the most forward rapidities. In order to estimate this effect, we have recomputed the $Z + jets$ background assuming that the calorimetry would not extend beyond $|\eta| = 4$, where the radiation dose for $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ is comparable to that at $\eta = 5$ for $10^{33} \text{ cm}^{-2}\text{s}^{-1}$. This loss of coverage will worsen the missing- E_t resolution of the detector, as shown in Fig. 3-29. Nevertheless, this figure shows that we are still dominated by real (detector independent) backgrounds and that therefore this decay mode remains viable at higher luminosity. This conclusion is valid because the signal is confined to missing- E_t greater than 300 GeV. The smaller missing- E_t signal that would result from the decay of a 300 GeV Higgs would be compromised at higher luminosity. However, as we have seen, such a particle would have already been discovered before the luminosity was raised.




FIG. 3-28. The distribution (solid histogram) in missing- E_t for the final state $Z(\rightarrow \mu^+\mu^-, e^+e^-) + \text{missing-}E_t$ including the effect of a Higgs boson of mass 800 GeV and the various backgrounds. The reconstructed Z is required to have $p_t > 250$ GeV and the events are rejected if they contain a jet with $E_t > 300$ GeV. The background shown as a dashed curve arises from $q\bar{q} \rightarrow ZZ$ (multiplied by 1.65 to account for the $gg \rightarrow ZZ$ process). The dot-dashed background arises from the final state $Z + jets$ where the missing- E_t is generated by calorimeter resolution or by losing energy out of the end of the detector. The dotted background arises from the final state $t\bar{t}$ where there is an e^+e^- (or $\mu^+\mu^-$) pair of mass $M_Z \pm 20$ GeV and the missing- E_t is due to neutrinos; the $Z + Q\bar{Q}$ background is negligible in this figure.




FIG. 3-29. The impact on the missing- E_t distribution for the $H \rightarrow ZZ \rightarrow \ell^+\ell^-\nu\bar{\nu}$ of reducing the coverage for the forward calorimeter to $|\eta| = 4$. The area between the solid curve and the dashed curve is the increased background as coverage is reduced from $|\eta| < 5$ to $|\eta| < 4$. The cuts are the same as those used for Fig. 3-28.

$H \rightarrow WW/ZZ \rightarrow \ell\nu/\ell\ell + 2 \text{ jets}$

The signal to noise ratio for the mixed hadronic and leptonic decays of the Higgs is much less favorable than that for the purely leptonic decays due to the large contributions from $W/Z + \text{jets}$ and $t\bar{t}$ production. Several phenomenological studies exist that suggest various means for reducing these backgrounds to a manageable level[26]. These analyses are very complex and involve a number of cuts which are very difficult to simulate accurately, due both to the reliability of the parton-shower Monte Carlo generators and to uncertainties in the modeling of the detector response. Here, we have not tried to present a complete signal to background analysis, but instead have concentrated on two general capabilities that are relevant for high mass symmetry breaking studies, and that should not be compromised in the detector design. The first involves reconstructing the mass and decay angular distribution for the $W/Z \rightarrow 2\text{-jet}$ decay, the second involves searching for the very forward jets produced by the WW/ZZ fusion process (Fig. 3-3 indicates that for large Higgs mass, this process dominates Higgs production at the SSC).

High p_t $W/Z \rightarrow 2\text{-jet}$ reconstruction

The first of these capabilities is the reconstruction of high p_t $W/Z \rightarrow 2\text{-jet}$ decays. For a Higgs mass in the region of 800 GeV or above, the interesting W/Z bosons have transverse momenta in excess of $M_{\text{Higgs}}/4$ (see Fig. 3-25). In this kinematic region, the W/Z decay products are highly boosted, resulting in two very narrow jets which are typically separated by 0.5 in (η, ϕ) space (see Fig. 3-19). These two jets will both normally lie inside of a jet definition cone of $R = 0.6$. This observation leads to the following strategy for reconstructing these jet pairs. First, find all of the jets using a large cone size of $R = 0.6$. Second, find all jets with a small cone size of $R = 0.15$ and match these narrow jets with the larger jet cones. To compute the mass of the system, all cells inside the larger cone are treated as massless 4-vectors and summed. The invariant mass of the sum is a good estimate of the W/Z mass[27]. The energies of the two narrow jets give a measurement of the decay angle:

$$\cos \theta^* \approx \frac{E_1 - E_2}{E_1 + E_2}.$$

In order to explore the viability of this approach, a sample of 1 TeV Higgs events with $H \rightarrow WW \rightarrow \ell\nu + 2 \text{ jets}$ has been generated using the HERWIG generator. After reconstruction of the events, we select those events in which the highest p_t jet has $p_t > 250$ GeV and $|\eta| < 2.5$. The mass of this jet in these events is shown in Fig. 3-30. A background sample is selected by applying the same criteria to a sample of $W + \text{jets}$ data containing ten times as many events as the signal sample (in the absence of any other cuts beyond the p_t requirement, the expected ratio of signal to background is 1 to 300). This sample is also shown in Fig. 3-30. It is apparent from this figure that the additional requirement that there are at least two narrow jets with $p_t > 25$ GeV inside the larger cone of $R = 0.6$ is a very powerful cut for rejecting backgrounds. The actual separation of the two peaks is somewhat sensitive to the size of the cone used in the jet definition—a larger cone size makes it easier for an ordinary QCD jet to acquire a mass of order M_W .

The mass resolution has been studied as a function of the calorimeter resolution and segmentation. The results are summarized in Table 3-7. Note that the fitted mean value of the W peak is also a strong function of $p_t(W)$, and increases by almost 10 GeV as $p_t(W)$ increases from 250 GeV to 750 GeV (this effect arises because, for the highly collimated jets at large $p_t(W)$ values, the finite cell size systematically increases the opening angle of the decay products, thereby increasing the observed mass). From this table, it appears that the performance of a calorimeter with 0.1 HAD1 segmentation is only slightly worse than that expected from a calorimeter with 0.05 HAD1 segmentation, with the exception of the highest bin in $p_t(W)$. As the segmentation is further increased, the mass resolution quickly degrades to an unacceptable level. It is worth noting that this conclusion depends strongly on including the effects of finite shower sizes. Simulations carried out at the particle level, ignoring shower spreading, show much greater deterioration of the mass resolution for segmentations in the range 0.1 to 0.2. Furthermore, the effects of extreme variations in the hadron calorimeter resolution (both stochastic and constant term), are rather minor, due to the strong dependence of the measurement on angular resolutions.




FIG. 3-30. The invariant mass distribution for the highest p_t jet in $H \rightarrow WW \rightarrow \ell\nu + 2 \text{ jets}$ events. A background contribution from $W + \text{jets}$ has been added to the plot to indicate the shape of the background (but the expected size of this background, in the absence of any other cuts, is thirty times larger than what is shown here). The jet mass has been computed using a cone of $R = 0.6$. The dotted curve indicates the size of the background before any requirement is made on the number of narrow jets identified inside the large cone. The dashed curve indicates the effect on the background of requiring at least two narrow jets with $p_t > 25$ GeV inside the larger cone. The solid histogram includes both signal and background after the requirement of two narrow jets inside the larger cone.

- (a) with an EM segmentation of 0.05 and a HAD1 segmentation of 0.05
- (b) with an EM segmentation of 0.05 and a HAD1 segmentation of 0.1
- (c) with an EM segmentation of 0.05 and a HAD1 segmentation of 0.2
- (d) with an EM segmentation of 0.05 and a HAD1 segmentation of 0.4

The ability of the calorimeter to reconstruct the individual narrow jets inside the larger jet definition cone of radius $R = 0.6$ is important for measuring $\cos\theta^*$. Note that after choosing a cone of $R = 0.15$ for the narrow jets, the jet definition no longer makes sense for cell sizes which are greater than or equal to 0.2. Once two narrow jets are found inside the larger cone, the energies of the two with the highest p_t are used to define $\cos\theta^*$. Using this definition, the reconstruction efficiency for $\cos\theta^*$ has been studied as a function of $p_t(W)$ and calorimeter segmentation. Recall that this variable is particularly interesting because, for large Higgs masses, the W/Z decay products are strongly longitudinally polarized (this effect represents the W/Z boson's memory that it gets its longitudinal component from the Higgs sector). This polarization leads to a decay distribution proportional to $\sin^2\theta^*$. The $W + \text{jets}$ background tends to produce a strongly forward/backward peaked distribution due to the soft bremsstrahlung spectrum of the second jet in these

Table 3-7

A summary of the mass resolution in GeV obtained for $W \rightarrow 2\text{-jet}$ decays as a function of $p_t(W)$, calorimeter resolution, and calorimeter segmentation (Δ is the cell size in (η, ϕ) space). The entries represent the sigma in GeV from a Gaussian fit to the distribution of the true mass minus the observed mass. For the upper three sets of values, only the HAD1 segmentation was changed, whereas for the bottom set, both the EM and HAD1 segmentation were changed. The top set is for a detector with a hadron calorimeter resolution described by a stochastic term of 40% and a constant term of 2%. The second and fourth sets are for the SDC baseline performance defined in Table 3-1. The third set is for a hadron calorimeter with a stochastic term of 100% and a constant term of 10%.

p_t range (GeV)	$\Delta = 0.05$	$\Delta = 0.1$	$\Delta = 0.2$	$\Delta = 0.4$
$250 < p_t(W) < 500$	4.00 ± 0.08	4.28 ± 0.08	5.92 ± 0.09	11.95 ± 0.22
$500 < p_t(W) < 750$	4.00 ± 0.10	4.54 ± 0.11	7.06 ± 0.15	12.96 ± 0.30
$750 < p_t(W)$	5.55 ± 0.33	6.20 ± 0.34	9.98 ± 0.57	24.4 ± 5.2
All	5.42 ± 0.09	5.90 ± 0.09	7.56 ± 0.11	15.31 ± 0.24
$250 < p_t(W) < 500$	4.99 ± 0.08	5.34 ± 0.08	6.63 ± 0.10	12.45 ± 0.22
$500 < p_t(W) < 750$	4.75 ± 0.11	5.31 ± 0.13	7.56 ± 0.16	13.26 ± 0.32
$750 < p_t(W)$	6.27 ± 0.30	6.92 ± 0.36	10.33 ± 0.65	17.7 ± 6.9
All	6.14 ± 0.09	6.63 ± 0.09	8.25 ± 0.11	15.84 ± 0.24
$250 < p_t(W) < 500$	5.82 ± 0.08	6.14 ± 0.09	7.34 ± 0.10	12.80 ± 0.22
$500 < p_t(W) < 750$	5.66 ± 0.12	6.17 ± 0.14	7.92 ± 0.17	13.14 ± 0.29
$750 < p_t(W)$	6.71 ± 0.38	7.43 ± 0.39	9.90 ± 0.58	24.5 ± 6.3
All	6.90 ± 0.09	7.38 ± 0.10	8.72 ± 0.11	15.77 ± 0.23
$250 < p_t(W) < 500$	—	5.88 ± 0.09	9.41 ± 0.13	18.83 ± 0.28
$500 < p_t(W) < 750$	—	6.52 ± 0.14	11.77 ± 0.19	21.85 ± 0.82
$750 < p_t(W)$	—	7.72 ± 0.41	16.35 ± 1.07	48.5 ± 15.1
All	—	7.38 ± 0.10	11.81 ± 0.15	21.28 ± 0.31

events. This is apparent in Fig. 3-31, which shows the observed $\cos\theta^*$ distribution from the Higgs signal and from the $W + \text{jets}$ background. The reconstruction efficiency for $\cos\theta^*$ in the signal events is fairly flat out to $|\cos\theta^*| < 0.8$ for $p_t(W) < 750$ GeV. Larger values of $\cos\theta^*$ correspond to asymmetric decays where the softer jet was outside the single reconstruction cone. Despite the decrease in pattern recognition information, the efficiency and resolution for $\cos\theta^*$ reconstruction for a HAD1 segmentation of 0.1 is found to be almost identical to that for a segmentation of 0.05.

Summarizing this discussion, the performance of the proposed SDC calorimeter has been studied for high p_t W/Z decays to 2-jet final states. The present design appears to have the minimum segmentation required for effective reconstruction of such final states.

Forward jet tagging

For large Higgs masses, the WW/ZZ fusion process becomes the dominant mechanism for producing the Higgs. This process, which has a unique kinematic signature that might be used to distinguish it from many backgrounds, initially produces W/Z boson pairs by radiating them from incoming quark lines. The WW/ZZ pairs then fuse to form the Higgs, leaving the parent quarks to continue on in the forward and backward directions (the typical transverse momenta of the radiated bosons are $\mathcal{O}(M_W)$). In principle, if a Higgs is discovered by other means, it should be possible to establish its couplings to W/Z bosons and t quarks by measuring the cross section for WW/ZZ fusion using jet tagging, and attributing the remainder of the observed production cross section to the gg fusion process which involves a t quark loop. This




FIG. 3-31. The distribution of $\cos\theta^*$ reconstructed for the highest p_t jet defined using a cone of $R = 0.6$. The decay angle $\cos\theta^*$ is defined using the energies of the two highest p_t narrow jets which fall inside the larger cone. The solid curve is for the $H \rightarrow WW \rightarrow \ell\nu + 2 \text{ jets}$ signal events, and the dotted curve is for the $W + \text{jets}$ background (the relative normalization is arbitrary).

measurement is very important in checking whether the observed Higgs couples as the Minimal Standard Model says it should, or whether there are likely to be other Higgs bosons which modify the expected couplings. In practice, the gg fusion process will produce jets from initial state radiation which can mimic the jet signature, and the jet tagging signature is somewhat ambiguous.

Two major issues need to be addressed in this context. The first is to determine the detector requirements for observation of the signal. To this end, a Monte Carlo sample of 1 TeV Higgs events has been generated using the HERWIG generator. This generator does not use the effective W approximation, and hence provides a better model of WW/ZZ fusion than other parton-shower Monte Carlos. The jets are defined in the conventional manner by clustering the observed transverse energy into cones in (η, ϕ) space with a radius of $R = 0.6$. In this analysis, the jet is required to have a minimum p_t in order to separate it from the soft scattering backgrounds. It is not yet known what the minimum p_t for reliable jet definition will be at the SSC; Table 3-8 includes results for two possible thresholds, the higher of which is almost certainly reliable. Furthermore, one of the defining features of the tag jets is their very high energy (they are initial state partons which have received small transverse momentum kicks). This naturally leads to a tag jet definition which uses the energy of the jet, after it has passed a minimum p_t requirement, to distinguish it from other jets.

This table suggests that efficient jet reconstruction for this process requires a fiducial coverage which extends to at least $\eta = 5$, while demonstrating that the effects of jet definition on the tagging efficiency in the baseline SDC forward calorimeter design are modest. It is also worth noting that the p_t of the tag jets is very low (the mean p_t for tag jets in the region $2.5 < |\eta| < 5$ is 100 GeV), and hence accurate measurements of their energy are not possible, due to the large smearing induced by intrinsic effects (fragmentation, etc.). This reduces the need for precision hadronic calorimetry in the forward region. Studies of the angular and energy resolutions have been performed as a function of the forward calorimeter segmentation (varying the cell sizes from 0.05 to 0.8 in $\Delta\eta \times \Delta\phi$) and performance (varying the performance from the baseline single-particle resolution given in Table 3-1 to that expected from a homogeneous 5 cm iron plate calorimeter with an EM resolution of $0.30/\sqrt{E} \oplus 0.02$ and a hadronic resolution of $0.80/\sqrt{E} \oplus 0.05$). In order to characterize the resolutions, several key jet observables were chosen (η , ϕ , p_t , and E), and an absolute resolution was defined using the generated partons. This definition includes the effects of fragmentation and hadronization, thereby providing a scale by which to measure the additional resolution effects induced by the detector. Figure 3-32 indicates that a $\Delta\eta \times \Delta\phi$ segmentation of 0.2 gives little degradation in angular or transverse momentum resolution when compared to 0.05 segmentation, and hence is well matched to the intrinsic resolution arising from jet definition effects. In fact, the p_t resolution is somewhat better for

Table 3-8

A summary of the acceptance for requiring one or two tag partons or jets in the forward region. The single tag case requires one parton or jet with $E > 3.0$ TeV, whereas the double tag case requires two partons or jets at opposite η with $E > 1.5$ TeV. The jets were reconstructed using cones with $R = 0.6$ in a forward calorimeter with cells of size $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$. The missing entries correspond to cases where the jets lie beyond the fiducial calorimeter coverage.

	Fiducial region	Parton $p_t > 25$ GeV	Parton $p_t > 50$ GeV	Jet $p_t > 25$ GeV	Jet $p_t > 50$ GeV
Single tag	$2.5 < \eta < 6$	0.40	0.32	—	—
	$2.5 < \eta < 5$	0.27	0.26	0.23	0.23
	$2.5 < \eta < 4$	0.076	0.076	0.068	0.068
Double tag	$2.5 < \eta < 6$	0.16	0.090	—	—
	$2.5 < \eta < 5$	0.10	0.072	0.078	0.052
	$2.5 < \eta < 4$	0.015	0.015	0.012	0.012

a segmentation of 0.2 than for 0.05 due to the large effect of the compartment threshold of $E_t > 100$ MeV for these low p_t jets (there are sixteen times as many cells which must pass the threshold cut for 0.05 than for 0.2). Furthermore, for the 0.2 segmentation, Fig. 3-33 indicates that the two different calorimeter resolutions are indistinguishable.




FIG. 3-32. The resolution for two tag jet observables in the $H \rightarrow WW \rightarrow \ell\nu + 2 \text{ jets}$ analysis as a function of the forward calorimeter segmentation. The solid curve is for a segmentation of 0.05 in both EM and HAD1. The dashed (dotted) are for 0.2 (0.4). The dot-dashed curve is for the extreme case of 0.8 (a) The ϕ resolution, defined to be $\phi(\text{true}) - \phi(\text{observed})$. (b) The p_t resolution, in percent, defined to be $(p_t(\text{true}) - p_t(\text{observed}))/p_t(\text{true})$.

The second major issue that must be addressed is to what extent the jet tagging signal just described is unique. The very forward region has many jets arising from initial state radiation which can then imitate the tag jets. The literature contains several studies at the parton level[28]. Here, the problem has been studied using the HERWIG generator. The results of this study are summarized in Table 3-9. These results indicate that there are large contributions from the backgrounds in the interesting regions, in contrast to the more optimistic assessments in the literature[28].

We have checked our results for the $t\bar{t}$ background against a recent complete $\mathcal{O}(\alpha_s^3)$ calculation of the $Q\bar{Q} + \text{jet}$ matrix element [29], and found reasonable agreement (see also Section 3.7.2). In order to




FIG. 3-33. The p_t resolution as defined in Fig. 3-32, for the two calorimeter designs defined in the text. The segmentation is fixed at 0.2 for both EM and HAD calorimetry. The solid curve is for a particle level simulation with no energy resolution effects, the dotted curve is for the SDC baseline resolution defined in Table 3-1. The dashed curve is for a calorimeter with an EM resolution of $0.30/\sqrt{E} \oplus 0.02$ and a hadronic resolution of $0.80/\sqrt{E} \oplus 0.05$ (significantly better than the baseline performance).

Table 3-9

A summary of the acceptance for requiring one- or two-tag jets for the signal and background processes of interest. No requirements on the event configuration in the central region have been made. The single tag case required one jet with $E > 3.0$ TeV, whereas the double tag case required two jets at opposite η with $E > 1.5$ TeV. The jets were reconstructed using cones with $R = 0.6$ in a forward calorimeter with cells of size $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ and were required to have $p_t > 50$ GeV and $2.5 < |\eta| < 5.0$.

	WW/ZZ fusion	gg fusion	$W + jets$ background	$t\bar{t}$ background
Single tag	0.23	0.082	0.11	0.035
Double tag	0.052	0.007	0.005	0.002

perform this check, we compare the fraction of $t\bar{t}$ events which contain additional jets in the matrix element calculation and in the parton-shower Monte Carlo. To approximate the matrix element calculation with our HERWIG event sample, we assume that the highest p_t additional jet is the one that should be compared with the matrix element (the remaining jets correspond to higher order effects). We find that the fraction of $t\bar{t}$ events with an additional jet of $p_t > 50$ GeV anywhere in the event is 50%, whereas the matrix element calculation gives a result of 44%. Confining our attention to the tag region of $|\eta| > 2.5$, we find 18% of the $t\bar{t}$ events have the highest p_t jet with $p_t > 50$ GeV in this region, whereas the matrix element predicts 13%. If we include all additional jets in the event with $p_t > 50$ GeV, the number rises to 23%, and if we also include the decay products from the t quarks, the fraction of events that have a jet with $p_t > 50$ GeV in the tag region reaches 38%.

This leads us to the conclusion that a forward calorimeter such as that in the SDC baseline design is vital for studying jet tagging physics. Although the efficiencies are low and the background rejections are not as large as one would like, there is clearly enough information present to untangle the two processes contributing to heavy Higgs production, and hence add a new dimension to such studies.

Conclusions

With the procedures discussed above for $W/Z \rightarrow 2\text{-jet}$ reconstruction and jet tagging, one can investigate the effects of simple cuts on the signal and the major backgrounds. For this purpose, we have used large event samples generated using HERWIG to study the $H \rightarrow WW \rightarrow \ell\nu + 2\text{ jets}$ final state. We choose the highest p_t jet and require it to satisfy $p_t > 250$ GeV. Then, two narrow jets with $R = 0.15$ and $p_t > 25$ GeV are required to be reconstructed inside the larger cone of radius $R = 0.6$, and the mass of the combined jet is required to lie in the W region ($65 < M_{\text{jet}} < 95$ GeV). These requirements have a 61% efficiency for the signal events, a 6% efficiency for $W + \text{jets}$ background, and a 2% efficiency for inclusive $t\bar{t}$ events (the $W + \text{jets}$ events were generated with the requirement that $p_t(W) > 250$ GeV). Next, the event is vetoed if there is a second jet found with $p_t > 25$ GeV in the central region ($|\eta| < 2.5$). This has a 71% efficiency for the signal, a 36% efficiency for $W + \text{jets}$ events, and a 7% efficiency for inclusive $t\bar{t}$ events. Finally, we require a single tag jet as defined in Table 3-9. Note that for the background, lack of Monte Carlo statistics forces us to assume that the rejection from the central cuts (the mass cut and the second jet veto) is independent of the rejection from the forward jet tagging, and therefore the total efficiency is the product of the central efficiency given above and the forward efficiency given in Table 3-9. For the signal events, the product gives a slightly smaller result than the full sequence of cuts (the 23% efficiency shown in Table 3-9 becomes 25% if the cuts are made in the correct sequence).

The result is a signal efficiency of 11%, and rejections of 500 against high p_t $W + \text{jets}$ events and 2×10^4 against inclusive $t\bar{t}$ events. When these factors are multiplied by the expected cross sections, the result is that both the $W + \text{jets}$ and inclusive $t\bar{t}$ backgrounds are reduced to roughly five times the size of the expected signal. The signal after cuts is expected to be about 300 events (including both electrons and muons) for a 1 TeV Higgs after one year at SSC design luminosity. Clearly, further reductions in the background rate must be sought in order to isolate the signal, but the large background reductions achieved using the simple cuts described here make this pursuit plausible, and emphasize the importance of reconstructing $W/Z \rightarrow 2\text{-jet}$ and tagging forward jets in the SDC detector.

3.2.4. SUSY extensions

The preceding Higgs discovery strategies relied on the properties of the Minimal Standard Model Higgs boson. In particular, the cross sections shown in Fig. 3-3 and the branching ratios shown in Fig. 3-4 are crucial ingredients. In this section, we consider a generalization of the Higgs sector, increasing the number of complex Higgs doublets from one to two. This leads to five observable Higgs states, which we denote as: h^0 (the lighter CP-even neutral Higgs), H^0 (the heavier CP-even neutral Higgs), A^0 (a CP-odd neutral Higgs), and H^\pm (the charged Higgs). The simplest and most popular form of this model is the Minimal Supersymmetric Standard Model (MSSM), which is described in more detail in Section 3.4.1.

In this model, the neutral component of one of the Higgs doublets is responsible for generating the mass of leptons and charge $-1/3$ quarks, while the other generates the masses of the charge $+2/3$ quarks. At the tree level, this model has two basic parameters, which are taken to be $\tan\beta$ (the ratio of the vacuum expectation values of the two Higgs doublets) and M_A . In the limit that M_A becomes very small (very large), the behavior of the heavier (lighter) Higgs H^0 (h^0) approaches that of the Minimal Standard Model Higgs. On the basis of renormalization group arguments, it is generally expected that $1 < \tan\beta < M_{\text{top}}/M_b$. There are some important mass relations which also exist in this model (for example, $M_{h^0} < M_Z$ and $M_{H^\pm} > M_W$). Radiative corrections to this model introduce several important changes to this picture (we assume that the mass scale of new supersymmetric particles such as squarks is large enough that they decouple from the calculations). The most significant of these is an expansion of the allowed mass regions for the lightest Higgs: $M_{h^0} < 110$ (140) GeV for $M_{\text{top}} = 150$ (200) GeV. This allows the lightest Higgs to move beyond the mass region accessible to LEP-II and into the mass region discussed in Section 3.2.1 and Section 3.2.2.

In order to assess more carefully what happens to the previous discovery scenarios, it is necessary to understand how the production cross sections and the branching ratios of the different Higgs bosons behave. The branching ratios are substantially modified, as shown in Fig. 3-34 and Fig. 3-35. The implications of

these changes (and additional changes in the production cross sections) have been analyzed in detail [30] for a total integrated data sample of 30 fb^{-1} (three SSC years at design luminosity). We summarize the conclusions in the following paragraphs.




FIG. 3-34. The branching ratios for the decay of the neutral Higgs bosons in the MSSM to the $\gamma\gamma$ final state, assuming $M_{\text{top}} = 150 \text{ GeV}$. The solid curve is for the h^0 , the dotted curve is for the H^0 , and the dashed curve is for the A^0 . The three different plots are for $\tan\beta = 0.5, 2$, and 20 .




FIG. 3-35. The branching ratios for the decay of the neutral Higgs bosons in the MSSM to the ZZ or ZZ^* final state, assuming $M_{\text{top}} = 150 \text{ GeV}$. The solid curve is for the h^0 , the dotted curve is for the H^0 , and the dashed curve is for the A^0 . The three different plots are for $\tan\beta = 0.5, 2$, and 20 .

The lighter Higgs (h^0) is either in the low mass (Section 3.2.1), or the intermediate mass (Section 3.2.2) region. For the lower mass region, the preferred strategy is to rely on the $h^0 \rightarrow \gamma\gamma$ decay mode, using the associated production processes $W + H$ and $t\bar{t} + H$. For large values of M_A , corresponding to the upper end of the allowed mass range for h^0 in Fig. 3-34, $BR(h^0 \rightarrow \gamma\gamma)$ approaches the Minimal Standard Model value, and discovery in this mode continues to be possible. For the intermediate mass region, where the preferred strategy is to search for $h^0 \rightarrow ZZ^*$, the only time that the h^0 mass is large enough for this branching ratio to become significant is for large values of $\tan\beta$ (see Fig. 3-35). In this case, the number of events expected is adequate for discovery.

The heavier CP-even Higgs H^0 is normally in the high mass region (Section 3.2.3). It has a reasonable branching ratio to ZZ only for values of $\tan\beta < 2$. This provides a window for discovery up to the threshold for $t\bar{t}$ production, at which point $H^0 \rightarrow t\bar{t}$ becomes the dominant decay. For larger $\tan\beta$ values, the H^0 decays predominantly to $b\bar{b}$ for H^0 masses below $t\bar{t}$ threshold. The overwhelming backgrounds, and

lack of a precise mass measurement, make both the $b\bar{b}$ and $t\bar{t}$ heavy quark decay modes for the H^0 very difficult to detect.

The charged Higgs is discussed in more detail in Section 3.3.3, but the situation can be briefly summarized. For all cases where the charged Higgs mass is at least 25 GeV less than the t -quark mass (corresponding to $BR(t \rightarrow H^\pm + b) > 1\%$), discovery appears possible. If the charged Higgs mass is higher than the t -quark mass, it has a small production cross section and it decays predominantly to $t\bar{b}$, making discovery very difficult.

The CP-odd Higgs A^0 is very difficult to find except for the case where $\tan\beta$ is small and M_A is below $t\bar{t}$ threshold. In this case, the decay mode $A^0 \rightarrow \gamma\gamma$ becomes observable, as suggested by Fig. 3-34. It is also potentially interesting to look for the decay $A^0 \rightarrow \tau^+\tau^-$. This final state is normally overwhelmed by other background sources, but there may be some regions of the parameter space where it becomes observable.

In conclusion, we summarize this discussion as a function of the location in the $(\tan\beta, M_A)$ parameter space, assuming an integrated luminosity of order 30 fb^{-1} . For small M_A , the h^0 is likely to be seen at LEP-II and the H^\pm is likely to be seen at the SSC. For moderate M_A and small $\tan\beta$, the h^0 is likely to be seen at LEP-II, the $H^0 \rightarrow ZZ$ is likely to be seen at the SSC, and the H^\pm might be seen in t quark decays at the SSC. For larger $\tan\beta$, none of the Higgs boson states may be easily observable. Finally, for large M_A , the h^0 is likely to be seen at the SSC instead of LEP-II, and the other Higgs bosons may be almost undetectable.

These conclusions are not strongly dependent on the particular detector under consideration. The significance of any narrow peak in ZZ^* or $\gamma\gamma$ will scale as the square root of the resolution, and hence only modest gains are possible. The major problem is that for large regions of the parameter space, the various Higgs bosons decay predominantly to heavy quarks, and thus become almost irretrievably buried by background. One possible technique is to use the $t\bar{t} + H$ production mechanism and look for events in which there are three or four b quarks tagged by the tracking system. Such complex final states have not yet been studied in detail.

3.2.5. Gauge boson pairs

We have already discussed the measurement of the ZZ and WW final states in the context of searches for Higgs bosons. Other final states of gauge boson pairs to be studied with the SDC detector include $Z\gamma$, WZ and $W\gamma$. In the Standard Model, the rates for these final states are well predicted. The WZ final state is particularly useful for testing our understanding of the quark distribution functions. It can be much more reliably estimated than the ZZ final state since it arises only from $q\bar{q}$ annihilation, and the QCD corrections are consequently much easier to calculate. They can be expected to be of the same order as those to $q\bar{q} \rightarrow ZZ$ [31] which are of order 20%. The rate for WZ production is approximately twice as large as that for ZZ .

The WZ final state can be reconstructed by the SDC detector using the final state $\ell\nu\ell\ell$, where ℓ is either an electron or muon. This mode has a combined branching ratio of 1%. A measurement of the missing transverse momentum together with the constraint of the W mass enables one to reconstruct the invariant mass of the WZ system, up to a possible twofold ambiguity. The only background once we require three isolated leptons, two of which reconstruct to a Z , arises from the $Zt\bar{t}$ final state. This final state was investigated as a possible background to $H \rightarrow ZZ \rightarrow \ell\nu\nu$ and found to be negligible compared to the intrinsic ZZ background. We assume that it will not be a problem for the WZ final state. Figure 3-36 shows the invariant mass distribution for W^+Z . Here events are selected that contain $e^+e^+e^-$. All three leptons are required to have $p_t > 20 \text{ GeV}$ and $|\eta| < 2.5$. Using the W mass as a constraint the longitudinal momentum of the neutrino is determined. There are in general two solutions, and if both are physical, they are both plotted. A comparison of the reconstructed and true WZ mass indicates that the SDC detector can reconstruct the invariant mass distribution of the pair very well. Perfect missing- E_t resolution was assumed in making this plot.

In technicolor models of weak interaction breaking, resonances may appear in the ZZ , WW , WZ or $Z\gamma$ channels. These resonances are analogous to the ρ and ω mesons of QCD. The masses of these resonances

are model dependent, but in order to detect them, assuming a sufficiently large number are produced, the SDC detector must be able to reconstruct the invariant mass of the boson pair with a resolution comparable to the width of the resonances. The width of the techni-rho (ρ_T^+) which decays to WZ is expected to be large (of order 25% of its mass). Mass resolution of this order can easily be achieved in the WZ channel since three of the decay products are measured completely and the neutrino's transverse energy can be inferred from the missing transverse energy.

The techni-omega has a much smaller natural width and can be detected via its decay to $Z\gamma$ [32]. Figure 3-37 shows the invariant mass distribution of the $Z\gamma$ where the Z is detected via its decay to lepton pairs. The small natural width of the techni-omega (of order 10 GeV) enables a peak to be clearly seen over the background which arises from $q\bar{q} \rightarrow Z\gamma$; the resonance is clearly visible at a mass of 1.46 TeV. There is an additional background from the final state $Z + \text{jets}$ where a jet fragments so that a single π^0 carries most of its energy. This background is a factor of 10 below that of $Z\gamma$, provided that the ratio of isolated π^0 's to jets at the same p_t (> 200 GeV in this case) is less than 0.001. The expected ratio is less than 5×10^{-4} .




FIG. 3-36. The distribution of invariant mass for the final state system of a W^+ and Z detected via the final state $e^+\nu e^+e^-$. One e^+e^- pair is required to have mass $M_Z \pm 10$ GeV. The leptons all have $|\eta| < 2.5$ and $p_t > 20$ GeV. In the events where there are two physical solutions for the neutrino's longitudinal momentum, both possibilities are plotted.




FIG. 3-37. The distribution of invariant mass for the final state system of a photon and an e^+e^- . The lepton pair is required to have mass $M_Z \pm 10$ GeV. The leptons have $|\eta| < 2.5$ and the photon has $|\eta| < 3.0$. The peak corresponds to the production and decay of the techni-omega particle of mass 1.46 TeV discussed in the text.

3.2.6. Strong breaking

The existence of fundamental scalar particles (Higgs bosons) with masses less than 1 TeV is only one of the possible manifestations of Electroweak Symmetry Breaking. An alternative scenario is that the scalar sector of the theory is strongly interacting (*i.e.*, the Higgs boson self-coupling is of order unity). These strong interactions will manifest themselves as modifications to the Standard Model scattering amplitudes for the longitudinal components of the W and Z bosons at large \hat{s} (this scattering is formally equivalent to the scattering of Goldstone bosons from the symmetry breaking sector [33]). In principle, it is possible to search for such effects in any gauge boson pair channel (ZZ , $W^\pm Z$, W^+W^- , or W^+W^+/W^-W^-).

If the strong interactions in the symmetry breaking sector produce resonances similar to those arising from QCD in $\pi\pi$ scattering (*e.g.*, as expected in Technicolor theories), one would expect large peaks in the W^+W^- , ZZ and $W^\pm Z$ invariant mass distributions, such as that shown in Fig. 3-36. A complementary approach involves a study of like-sign W production [34]. In particular, in a pp machine, W^+W^+ production

is approximately a factor of three larger than W^-W^- production. This is expected to be a non-resonant channel, and is complementary in that one expects that if the resonant signals are suppressed, this non-resonant channel is enhanced, and vice versa. It has an additional advantage in that the lowest order $q\bar{q}$ and gg diagrams which produce large numbers of transversely polarized background events in the channels W^+W^- , WZ , and ZZ do not exist for the W^+W^+ channel. A number of phenomenological studies exist[35] that suggest it is possible to isolate the very small expected signal after a number of complex cuts.




FIG. 3-38. The lepton transverse momentum distribution expected for the W^+W^+ signal and backgrounds (there are two entries per event, one for each lepton). The signal is shown as the solid histogram, the dashed histogram is for $t\bar{t} \rightarrow W^+W^- + X$, and the dot-dashed histogram is for $t\bar{t} \rightarrow W^+b\bar{b} + X \rightarrow \ell^+\ell^+ + X$. Note that the curves for these two backgrounds sources lie on top of each other. The dotted histogram is for the direct production of opposite-sign $q\bar{q} \rightarrow W^+W^-$ background.




FIG. 3-39. The curvature distribution for a sample of events consisting of lepton tracks with $p_t = 500$ GeV superimposed on a minimum bias background corresponding to a luminosity of $3 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$. The curves represent Gaussian fits to the data, and demonstrate the absence of any non-Gaussian tails towards small values of the curvature (the relevant aspect for charge mis-measurement studies). (a) The distribution for $p_t = 500$ GeV electrons. (b) The distribution for $p_t = 500$ GeV muons.

There are two major sources of experimental backgrounds. The first arises from opposite-sign W pairs (produced either directly by $q\bar{q} \rightarrow W^+W^-$, or indirectly by $gg \rightarrow t\bar{t} \rightarrow W^+W^- + X$) in which the charge of one of the relevant tracks is mis-measured. Note that the direct production process is indistinguishable from the signal if the charges are mis-measured, whereas in the case of the $t\bar{t}$ process, there are other

potential means to distinguish signal and background (*e.g.*, extra jet activity in the event). The second background arises from like-sign W pairs coming from $gg \rightarrow t\bar{t} \rightarrow W^+b\bar{b} + X \rightarrow \ell^+\ell^+ + X$. Figure 3-38 shows the expected lepton transverse momentum distribution for the signal and these backgrounds, without any p_t requirements (both leptons are required to have $|\eta| < 2.5$). The signal is derived from a calculation by Chanowitz (Ref. 34), in which the lowest partial wave is allowed to increase linearly until it saturates at the unitarity bound. This provides a signal somewhere in the middle of the range of allowed models. The backgrounds have been computed using PAPAGENO. From this figure, we infer that the fraction of leptons with mis-measured charge is required to be less than 10^{-5} at $p_t < 100$ GeV, and less than 10^{-3} at $p_t = 500$ GeV in order to completely suppress the opposite-sign backgrounds. For the like-sign background, the charges are correct but one of the two leptons is not isolated. Studies carried out in Section 3.2.2 indicated that rejections of about 1000 were obtainable for leptons with $p_t > 100$ GeV by requiring the excess E_t in a cone of radius $R = 0.3$ around the lepton to be less than 5 GeV. Furthermore, the $t\bar{t}$ events will have additional jet and lepton activity in the detector which can be used to veto them. Although we have not performed the necessary simulations to confirm that the required rejections are feasible, it appears plausible that an additional factor of 100 (on top of the factor of 1000 obtained from local isolation) could be obtained [35]. Figure 3-38 indicates that this should be sufficient to see the signal.

Due to the vital role that charge separation plays in this analysis, the ability of the SDC detector to separate charges has been studied in more detail. The parametrized resolution of the tracking system is adequate to achieve the stated goals for $|\eta| < 2.0$, and becomes more marginal beyond this value (Fig. 3-1 shows that at 1 TeV, the resolution is 15% for $|\eta| < 1.6$), but complete tracking simulations are needed, including all of the underlying tracks in the event, in order to search for non-Gaussian tails on the momentum resolution. Figure 3-39 presents the result of such a simulation, performed at a luminosity of $3 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$. For this analysis, a sample of lepton tracks with $p_t = 500$ GeV has been reconstructed, including the appropriate background events, and also including all of the effects of hit generation, pattern recognition, and fitting. This simulation covers the region $|\eta| < 1.6$ only because the present simulations of the intermediate tracking system are still under development. Due to the difficulty of performing such simulations, we have chosen to generate the event sample at large p_t , where sign measurement errors are most likely. Several effects might appear. First, non-Gaussian errors might arise from confusion caused by the additional hits due to nearby soft tracks. Second, in the case of electrons, the presence of bremsstrahlung from the material in the tracking volume could occasionally cause a large momentum mis-measurement in the wrong direction. Figure 3-39 plots the appropriate variable (the track curvature, defined to be $\rho = 0.03B/p_t$, where B is in tesla and p_t is in TeV) for samples of 900 electron and muon tracks. The tracks have been selected to be of high quality by requiring a small χ^2 from the fitting procedure. The reconstruction and fitting procedure, including the quality cuts, has an efficiency of roughly 80% for both electrons and muons. Only a single muon event was observed to have the wrong sign for the curvature, and there are no other events even close to the wrong sign. This leads us to set an upper limit, based on our current somewhat immature track reconstruction codes, of 10^{-3} for the wrong sign charge rejection at $p_t = 500$ GeV. Simulations with larger event samples will be required to extend this limit to lower values of the transverse momentum where the charge separation is expected to be significantly better.

3.3. Physics of the t Quark

The t quark is one of the few ingredients of the Standard Model that has not been directly observed. The discovery of this particle and studies of its properties will stringently test the Standard Model. Furthermore, since the t quark production cross section at the SSC is expected to be large (~ 10 nb for a t -quark mass of 150 GeV), it is potentially a serious background to much rarer processes such as the production of a Higgs boson followed by its decay into a W^+W^- final state (see Section 3.2.3). Extensions to the Standard Model involving a more complex Higgs sector may also be most readily observed in the decays of t quarks. Finally, recent calculations have indicated that the associated production of Higgs particles with t quarks is quite large at the SSC. The tagging of t quarks can therefore be an effective signature for a Higgs particle search (see Section 3.2.1). For these reasons, the SDC detector has as one of its design goals the ability to efficiently detect and study t quark production and decay.

Direct searches for the t quark have been made at e^+e^- and $p\bar{p}$ colliders, and these searches have ruled out a t quark with mass less than 45 and 91 GeV at 95% CL, where the first limit is independent of the t quark decay mode [36] while the higher limit is obtained when one assumes that it decays to a $W^\pm b$ final state as predicted by the Standard Model [38]. Recent hadron collider measurements of the W width [37] also constrain the mass of the t quark to be above 55 GeV at 95% CL, independent of the t quark decay mode. Combined fits, including the recent precision measurements of the Z mass, width and decay asymmetries, as well as deep-inelastic results and the recent precise measurements of the W mass from hadron colliders, predict an allowed region for the t -quark mass. It is restricted to lie in the range $90 < M_{\text{top}} < 200$ GeV at 95% CL, independent of the t quark decay modes. This limit is valid, within the context of the Minimal Standard Model, for $50 < M_{\text{Higgs}} < 1000$ GeV [39]. The CDF and D0 experiments at Fermilab may very well discover the t quark within the next few years. In this case, the SSC will be a copious source of t quarks, and detailed study of t quark pair production will be possible. If the t quark remains undetected by the time the SSC is commissioned, the search for this particle will be one of the first challenges facing the SDC experiment.

The production cross section for t quarks, shown in Fig. 3-40 as a function of the t -quark mass, has been calculated to next-to-leading order, and has an uncertainty of $\sim 25\%$ due to structure function and Q^2 -scale uncertainties. The dominant production mechanism at the SSC is gluon fusion that produces t quark pairs with a t quark p_t spectrum peaking around half the mass of the t quark. Since the Standard Model predicts the t quark to decay almost completely to $W^\pm b$ final states, $t\bar{t}$ production will result in events with typically two high- p_t W^\pm bosons and two b quarks. The SDC detector is ideally suited to identify such events by tagging one or both W^\pm bosons through their leptonic decay modes, and tagging b quarks with the silicon tracking system.




FIG. 3-40. The production cross section for heavy t quarks. The band corresponds to the envelope of the smallest and largest rates obtained by varying the factorization scale between 0.5 and 2 times the quark mass, and by using the following sets of structure functions [40]: DFLM160, DFLM360, HMRSB, and MTB.

In the following sections, we will review the techniques we have developed to study t quark decays with the SDC detector. We first outline several methods for detecting the t quark and measuring its mass. We then outline two methods that can be used to search for a charged Higgs particle produced in the decay of a heavy t quark.

We note below that the measurement of the t quark branching fractions can be used to search for a charged Higgs particle. However, this also provides a means of indirectly detecting the coupling of the t quark to other final states not predicted by the Standard Model. The accuracy with which the branching fractions can be measured depends on our understanding of the systematic uncertainties associated with the backgrounds in each channel. We therefore will not discuss this topic in full generality, and instead

focus on the specific case of the charged Higgs as an example of how the SDC detector can perform such searches.

The analyses described below have used the ISAJET event generator to model $t\bar{t}$ production and decay. We have modelled the detector response using the parametrizations described in Section 3.1.1 which incorporate effects resulting from the finite calorimeter segmentation, as well as the magnetic field and the calorimeter non-linearities.

3.3.1. Detection and mass measurement for the t quark

We have considered three methods for detecting t quarks and measuring the t -quark mass. The first technique involves detecting the final states produced when one t quark decays to $e\nu_e + b$ while the other decays to $\mu\nu_\mu + b$. The second method involves searching for "sequential" t quark decays, where a t quark decays semileptonically and the decay of the associated b quark results in a non-isolated muon. The third method uses the events resulting from the semileptonic decay of one t quark into either $e\nu_e + b$ or $\mu\nu_\mu + b$ and the hadronic decay of the other t quark into three jets.

Direct t quark decays to dilepton final states

The cleanest signature for t quark production is expected to be the isolated high- p_t electron and muon that result from the semileptonic decays of both t quarks. The requirement of two different lepton species avoids backgrounds from Z^0 and Drell-Yan production, and the rate of $e^\pm\mu^\mp$ events is twice the rate of either e^+e^- or $\mu^+\mu^-$ final states. The most effective dilepton search would include these latter two final states (as demonstrated recently by the CDF collaboration[38]) but the detection efficiency will be reduced relative to the $e^\pm\mu^\mp$ mode because of the need to reject the backgrounds mentioned above. We therefore focus our discussion on the $e^\pm\mu^\mp$ final states.

Candidate events are selected by requiring isolated electron and muon candidates of opposite charge, with $p_t > 20$ GeV and $|\eta| < 2.5$. This selection is expected to yield 10^6 events per SSC year for $M_{\text{top}} = 150$ GeV, and should be essentially free of background. Thus the analysis becomes a counting experiment, and we can use the observed number of candidate events to estimate the t quark production cross section. Since the cross section falls as a function of t -quark mass, we can use the observed $t\bar{t}$ cross section to estimate the mass, assuming the Standard Model t quark branching ratios into semileptonic final states. The uncertainty in this t mass measurement is dominated by the uncertainty in the predicted cross section, shown in Fig. 3-40, and is estimated to be 10-15 GeV.

The uncertainty in the t -quark mass determined in this way is quite large, and depends on the assumption of the Standard Model branching ratios. We therefore consider this analysis to be a clean method for discovering the t quark or confirming its existence, but it is not competitive with other methods for measuring the mass.

The sequential $e\mu$ method for t -quark mass determination

A second method for measuring the t -quark mass relies on events with one isolated electron (from t decay) and one non-isolated muon of opposite sign (from the b -decay product of the same t quark) [41]. We have studied this technique for t -quark masses of 150 and 250 GeV. The inclusive electron trigger with $p_t > 40$ GeV is ideal for obtaining a sample of $e^\pm\mu^\mp$ events for this analysis. In addition, we require the electron to be isolated by selecting events with less than 4 GeV of excess transverse energy in a cone of radius $R = 0.2$ around the electron direction (where $R^2 = (\Delta\eta)^2 + (\Delta\phi)^2$).

Muon candidates from the semileptonic b decay are selected by requiring them to have $p_t > 20$ GeV. In conjunction with the electron requirement, this is expected to yield a sample of clean $e^\pm\mu^\mp$ events. Non-isolated muons are selected by requiring at least 20 GeV of excess transverse energy in a cone of radius $R = 0.4$ around the muon direction. A muon of opposite sign to the electron can also arise from a charm decay generated in the cascade decay of the other t quark in the event. This background muon will generally lie in the opposite hemisphere from the electron, as shown in Fig. 3-41. There is a clear separation

between the sequential muons (*i.e.*, those muons coming from the same t quark as the electron) and the muons from the charm background. We reduce the charm background by requiring that the azimuthal distance between the electron and the muon, $\Delta\phi$, be less than 80° . Finally, we require $p_t(e\mu) > 100$ GeV to increase the sensitivity to the t -quark mass. For the case of $M_{\text{top}} = 250$ GeV, this requirement is modified to be $p_t(e\mu) > 120$ GeV.

Figure 3-42 shows the invariant mass of the $e\mu$ pair, $M(e\mu)$, for t -quark masses of 150 GeV and 180 GeV with the above cuts. The events from the heavier t -quark peak at a higher $M(e\mu)$. Figure 3-43 shows a similar plot for t -quark masses of 220 and 250 GeV. The sensitivity to the $p_t(e\mu)$ cut can be seen in Fig. 3-44 as a function of M_{top} . It shows the mean invariant mass of the $e\mu$ pair as a function of the t quark mass for several values of the $p_t(e\mu)$ cut. The mean $M(e\mu)$ has an approximately linear dependence on the t -quark mass and the sensitivity (the slope) increases with the $p_t(e\mu)$ cut. Transverse momentum cuts of 100 GeV and of 120 GeV for the t -quark masses of 150 and 250 GeV, respectively, give adequate sensitivity while retaining sufficient statistics for a good mass determination. In one year of running at nominal SSC luminosity, we expect 70,000 and 17,000 events of this type for the two mass values, providing a measurement of the t -quark mass with a statistical uncertainty of 0.5 (0.8) GeV for t -quark masses of 150 (250) GeV.




FIG. 3-41. The distance in $\eta - \phi$ space between the electron and muon candidates for events selected as explained in the text. The two peaks, from left to right, correspond to non-isolated muons from b and c decays, respectively. This distribution is for the case of $M_{\text{top}} = 150$ GeV.




FIG. 3-42. The invariant mass distribution for the $e\mu$ pair, $M(e\mu)$, for two different t -quark masses. The lefthand scale and the leftmost curve are for $M_{\text{top}} = 150$ GeV, while the righthand scale and curve are for $M_{\text{top}} = 180$ GeV. The cut $p_t(e\mu) > 100$ GeV has been used. The superimposed curves represent Gaussian fits.

In this analysis, the backgrounds from other processes are very small. We have considered WW , $Z \rightarrow \tau\tau$, and $W + b\bar{b}$ production. Only the last process is a non-negligible source of isolated electrons and non-isolated muons. We have estimated this background using the ISAJET Monte Carlo program and find that this process contributes a 3% background to the $M(e\mu)$ plot for a t -quark mass of 250 GeV and a 0.7% background for a t -quark mass of 150 GeV.

The systematic uncertainty on the t -quark mass using this method is dominated by uncertainties in the physics inputs. The first is the incomplete knowledge of the b -quark fragmentation function, which affects the muon momentum distribution. We use the Peterson fragmentation parametrization for heavy quarks with the value for the ϵ parameter measured by ALEPH [42]. Variations of ϵ by one standard deviation result in 1.5 GeV and 3.0 GeV variations in the measured t -quark mass for $M_{\text{top}} = 150$ and 250 GeV, respectively. The second uncertainty arises from the imprecise knowledge of the t -quark p_t distribution. To




FIG. 3-43. The invariant mass distribution for the $e\mu$ pair, $M(e\mu)$, for two different t -quark masses. The leftmost scale and the leftmost curve are for $M_{\text{top}} = 220$ GeV, while the righthand scale and curve are for $M_{\text{top}} = 250$ GeV. The cut $p_t(e\mu) > 120$ GeV has been used. The superimposed curves represent Gaussian fits.




FIG. 3-44. The mean invariant mass of the $e\mu$ pair as a function of the t -quark mass for $p_t(e\mu) > 60$ GeV (lower curve), $p_t(e\mu) > 100$ GeV, and $p_t(e\mu) > 140$ GeV (upper curve). The data are derived from Gaussian fits to distributions such as those shown in Fig. 3-42.

illustrate this point, the dependence of $M(e\mu)$ on the t quark p_t for the case of a t quark with a mass of 150 GeV is shown in Fig. 3-45. In addition, we expect different t -quark p_t spectra for other t -quark production processes, such as $W + t$ production or $t\bar{t}$ -pair production from gluon splitting. The dependence of $M(e\mu)$ on the t -quark p_t spectra has been studied by reducing the initial state radiation generated by ISAJET to obtain a softer p_t distribution for the t quark. From this variation, we derive uncertainties on the t -quark mass of ± 1.9 GeV and ± 2.6 GeV for $M_{\text{top}} = 150$ and 250 GeV, respectively.




FIG. 3-45. The invariant mass of the $e\mu$ pair as a function of the t -quark p_t for $p_t(e\mu) > 100$ GeV. The data are derived from Gaussian fits to distributions such as those shown in Fig. 3-42, for different t -quark p_t bins.

Both of these systematic uncertainties may be reduced after detailed studies of data from the SSC. Adding all the uncertainties in quadrature, we expect to determine the t -quark mass at 150 GeV with an uncertainty of ± 0.5 (stat) ± 2.4 (syst) GeV and at 250 GeV with an uncertainty of ± 0.8 (stat) ± 3.9 (syst) GeV after a run of one year at the SSC design luminosity.

3.3.2. Decays of $t\bar{t}$ to lepton + jets final states

The t -quark mass can be measured directly by reconstructing the three jets coming from its hadronic decay, $t \rightarrow bu\bar{d}$ or $bc\bar{s}$. In $pp \rightarrow t\bar{t}$ production, one can trigger on a lepton from the decay of one t quark into the modes $b\nu_e$ or $b\nu_\mu$, and then examine the three jets from the hadronic decay of the other t quark. These events typically have a high- p_t charged lepton that is well-isolated from other energy flow in the event. The backgrounds to this signature are leptons from b quark decay, which tend to be non-isolated and to have lower p_t , and leptons from the decay of W bosons produced inclusively. These backgrounds can be reduced to a negligible level with the cuts described below.

We have studied the capability of the SDC detector to measure the t -quark mass in this channel for the two cases $M_{t\text{op}} = 150$ and 250 GeV. The identical analysis procedure was used to study the t quark decay via a charged Higgs boson, which is described in a subsequent section. The ISAJET 6.36 Monte-Carlo was used to generate samples of $t\bar{t}$ events at the two masses. In both cases, the generated luminosity corresponded to a small fraction of an SSC year (about 3% for the 150 GeV case, and 8% for the 250 GeV case). The subsequent histograms have been scaled up to contain the correct numbers of events, but the statistical fluctuations visible in the plots are much larger than one would expect for the numbers of events they are claimed to contain. Additional underlying events, corresponding to design luminosity, were not added. The events were simulated with the calorimeter model described in Section 3.1.1, which includes the effects of finite calorimeter segmentation and non-linear response as well as the magnetic field. Jets were reconstructed using the clustering algorithm described in Section 3.1.1 with a cone size $R = 0.4$ in the region $|\eta| < 2.5$. This small cone size was chosen in order to minimize losses in efficiency in these complex multi-jet events, while still maintaining a reasonable energy measurement.

For this analysis, we select events in which the t and \bar{t} are produced recoiling against each other with high transverse momentum. The event rates and efficiencies for the cuts described below are listed in Table 3-10. We require an electron or muon with $p_t > 40$ GeV and $|\eta| < 2.5$, consistent with the expected inclusive lepton trigger threshold. The leptons are required to be isolated by demanding that the additional E_t in a cone of radius $R = 0.4$ about the lepton is less than 25% of the lepton p_t . The lepton efficiency for these cuts is higher for the 250 GeV case, due to the more central rapidity distribution and harder p_t spectrum.

To select events with a hadronic t quark decay, we require that at least three jets be reconstructed, each with measured $p_t > 30$ GeV in the opposite hemisphere from the lepton ($\Delta\phi > 90^\circ$). We note that this p_t cut is made before corrections to the jet energy scale for detector effects have been applied; this data is referred to as "uncorrected" in the subsequent discussion. One of these jets must be a b jet, within $|\eta| < 2$, which is tagged with a secondary vertex (see the tracking section for more details). Based on our studies, we estimate the efficiency for the b tag to be 30% for b jets with $p_t > 30$ GeV and inside the acceptance discussed above, with a background rejection of order 100 to 1 against non- b jets. In principle, one could also use semileptonic b tags [44], but since these involve neutrinos they could bias the t -quark mass measurement.

In order to reduce combinatoric background, we require the transverse momentum of the three-jet system to be large, which collimates the three jets, and separates them from other jets in the event. The p_t distribution for three-jet combinations opposite the lepton is plotted in Fig. 3-46. If there is more than one three-jet combination which passes the selection criteria in a given event, then multiple entries are plotted for that event. For the 150 and 250 GeV cases, we choose the minimum three-jet p_t to be 200 and 300 GeV, respectively, resulting in approximately 160,000 and 40,000 three-jet combinations (t candidates) per nominal SSC year.

The invariant mass distributions for the two jets, not including the b jet, are plotted in Fig. 3-47 (a) and Fig. 3-47 (b), respectively, and the three-jet invariant mass distributions are shown in Fig. 3-48 for the 150 GeV and 250 GeV cases. Large invariant mass peaks for the W and t are evident and the combinatoric background is relatively small. Note that in these plots the jet energies have not been corrected for detector and jet definition effects, and therefore, the peaks appear at lower invariant mass than the actual particle masses (73.9 GeV for the W , and 134.1 GeV for the t , for the $M_{t\text{op}} = 150$ GeV case).

Table 3-10
Summary of event samples and efficiencies for the event selection for $M_{\text{top}} = 150$ and 250 GeV.

	$M_{\text{top}} = 150 \text{ GeV}$	$M_{\text{top}} = 250 \text{ GeV}$
$\sigma(t\bar{t})$	12 nb	1.5nb
$N_{\bar{t}}$	1.2×10^8	1.5×10^7
Branching ratio	8/27	8/27
Lepton, geometric	0.43	0.56
Lepton id, isolation	0.85	0.85
b jet, geometric	0.51	0.71
b jet tag	0.30	0.30
$N(p_t(3\text{-jet})) > 200 \text{ GeV}$	158,000	—
$N(p_t(3\text{-jet})) > 300 \text{ GeV}$	—	39,400

In general, the jet energies need to be corrected for three effects: (i) non-linearity and other losses (cracks, neutrinos, leakage) which cause the energy deposited in the jet definition cone to be underestimated; (ii) energy from the jet that falls outside the jet definition cone (either produced at large angles or swept away by the magnetic field); and (iii) energy entering the cone from the underlying event or background events. In Fig. 3-49, we plot the fractional difference between the measured jet p_t and the "true" jet p_t , as a function of the "true" jet p_t , where the true value ($p_t(\text{true})$) is defined as the transverse momentum sum of particles (excluding neutrinos) inside a cone of $R = 0.4$ about the direction of the parent. In the present simulation, there are no crack effects included, and hence the major effects entering the correction factor are sweeping of charged tracks by the magnetic field and calorimeter non-linearities. The dotted curve shows an average fractional loss which will be corrected in subsequent plots.




FIG. 3-46. The three-jet transverse momentum distribution for three-jet combinations (t candidates) opposite the lepton in $t\bar{t} \rightarrow WWb\bar{b}$ events. The t quark mass was taken to be 150 GeV.

To further reduce the background to the two-jet invariant mass distribution, we require that the corrected three-jet invariant mass be in the range $135 < M(3\text{-jet}) < 165 \text{ GeV}$ ($225 < M(3\text{-jet}) < 275 \text{ GeV}$ for the 250 GeV case). The two-jet mass distributions after this cut are shown in Fig. 3-50 (a) and (b). After correcting for the fractional loss displayed in Fig. 3-49, the W mass peak appears at 80.5 GeV with a width of 7.5 GeV, to be compared with the generated W mass of 80.0 GeV.




FIG. 3-47. The observed (uncorrected) two-jet invariant mass distribution using the cuts described in text for (a) $M_{top} = 150$ GeV and (b) 250 GeV.




FIG. 3-48. The observed (uncorrected) three-jet invariant mass distribution for (a) $M_{top} = 150$ GeV and (b) $M_{top} = 250$ GeV.

Likewise, if we require the two-jet invariant mass to be in the range 65 to 95 GeV, the three-jet invariant mass distributions of Fig. 3-51 (a) and (b) show almost no background. The measured t invariant masses are 147.9 GeV and 243.6 GeV for the 150 (250) GeV true mass. For the lower t -quark mass, the peak has a width of roughly 9 GeV, whereas the higher t -quark mass peak has a width of about 14 GeV.

The t -quark mass can be inferred from the three-jet invariant mass distribution with a statistical accuracy of ~ 0.04 GeV after one year of operation at standard luminosity. The observed cross section is large enough that this measurement can be reliably made with a small fraction of this integrated luminosity, and this statistical uncertainty will simply scale as $\sqrt{N_{cand}}$, where N_{cand} is the number of observed candidates. We expect the t -quark mass uncertainty to be dominated by systematic uncertainties, with the largest contribution coming from the uncertainty in the calorimeter energy scale. However, the W signal in the dijet channel in these final states can be used to calibrate the calorimeter mass scale. We conservatively estimate the remaining systematic uncertainty to be ~ 3 GeV, after taking into account possible differences in energy response for light quark and heavy quark jets. Another method to reduce the uncertainty would be to use the tracking system to provide a track-by-track correction for non-linearity; in CDF this was shown to reduce the systematic uncertainty in the jet energy scale calibration and improve the jet energy resolution by 10–15% [43].




FIG. 3-49. The fractional loss of energy in a jet due to detector effects. The definition used here is: $(p_t(\text{observed}) - p_t(\text{true}))/p_t(\text{true})$, where $p_t(\text{true})$ is the transverse momentum sum of particles (excluding neutrinos) inside the $R = 0.4$ jet cone, and $p_t(\text{observed})$ is the measured energy in the jet cone, including all detector effects. The dotted curve shows the correction factor applied in subsequent plots.




FIG. 3-50. The corrected two-jet invariant mass distributions after requiring that the three-jet invariant mass be consistent with the t mass. (a) $M_{\text{top}} = 150$ GeV, requiring $135 < M(3\text{-jet}) < 165$ GeV. (b) $M_{\text{top}} = 250$ GeV, requiring $225 < M(3\text{-jet}) < 275$ GeV.

3.3.3. Decays of the t quark to charged Higgs bosons

Introduction to charged Higgs bosons

One of the most attractive extensions of the standard Higgs sector contains two Higgs doublets and consequently both charged and neutral physical Higgs bosons (see Section 3.2.4) [45]. If the charged Higgs boson is lighter than the t quark, then the branching ratio for the decay $t \rightarrow H^+ b$ could be comparable to that for $t \rightarrow W^+ b$. These branching fractions depend on the couplings of the two Higgs doublets to the quarks and leptons. There are two possible models normally considered for these couplings consistent with the absence of flavor-changing neutral currents. In the most popular model, the neutral component of one of the doublets is responsible for generating the mass of leptons and charge $-1/3$ quarks while the other generates the mass of charge $+2/3$ quarks. This is the model predicted by minimal supersymmetry and will be the one considered here. The couplings of the charged Higgs bosons to fermions are entirely determined by the quark/lepton masses and by $\tan\beta = v_2/v_1$, where v_1 (v_2) is the vacuum expectation




FIG. 3-51. The corrected three-jet invariant mass distribution, for (a) $M_{\text{top}} = 150$ GeV and (b) $M_{\text{top}} = 250$ GeV. In this plot the two-jet invariant mass is required to be consistent with the W mass $65 < M(2\text{-jet}) < 95$ GeV).

value of the Higgs field which couples to the down (up) type fermions. Therefore, $\tan\beta$ determines the branching fractions for $t \rightarrow bH^+$, $H^+ \rightarrow \tau\nu$, and $H^+ \rightarrow c\bar{s}$.




FIG. 3-52. The branching ratios for the reaction $t \rightarrow H^+b$ as a function of $\tan\beta$. We have assumed various values of M_{top} and M_{H^\pm} as labeled.




FIG. 3-53. The branching ratios for the reactions $H^+ \rightarrow \tau\nu$, $c\bar{s}$, $c\bar{b}$ as a function of $\tan\beta$.

The predicted branching ratios for $t \rightarrow H^+b$ as a function of the parameter $\tan\beta$ are shown in Fig. 3-52 for several values of M_{top} and M_{H^\pm} . The various H^\pm branching ratios, which are essentially independent of M_{top} and M_{H^\pm} , are shown in Fig. 3-53 again as a function of the parameter $\tan\beta$. Very small values of $\tan\beta$ ($\lesssim 0.2$) would place the $H^+ \rightarrow t\bar{b}$ coupling in a non-perturbative regime since M_{top} is large. One particularly important feature of Fig. 3-52 and Fig. 3-53 is that the branching ratios for $t \rightarrow H^\pm b$ and $H^\pm \rightarrow \tau\nu$ tend to be anti-correlated for $\tan\beta$ below 10. For example, as $\tan\beta$ approaches 10, the t branching fraction to H^\pm is approximately 1%, while the H^\pm decays almost entirely into $\tau\nu$. Note also, that for small $\tan\beta$ the $\tau\nu$ branching ratio is quite small, and the $c\bar{s}$ mode will provide the best hope for H^\pm detection.




FIG. 3-54. The ratio of the measured single track momentum to the energy within a cone of radius $R \equiv \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} = 0.4$ for pions from τ decays. The effects of hadron shower size have been included in this simulation.

Event selection for charged Higgs searches

We have investigated two methods [46] for H^\pm detection in $t\bar{t}$ events using several different values of M_{top} and M_{H^\pm} . The first involves a search for an excess of τ leptons. This technique is most effective when the branching ratio for $H^+ \rightarrow \tau\nu$ is large. The other method is to reconstruct the hadronic decays $H^+ \rightarrow c\bar{s}$. It is useful for smaller values of $\tan\beta$ where $t \rightarrow H^+b$ and $H^+ \rightarrow c\bar{s}$ are both large.

In each case, events are triggered by requiring one t quark to decay via $t \rightarrow bW \rightarrow b\ell\nu$ yielding an electron or muon (ℓ) with $p_t > 40$ GeV and $|\eta| < 2.5$. We require these leptons to be isolated from energy flow in the rest of the event by demanding that the energy within a cone of radius $R = 0.4$ about the lepton be less than 25% of the lepton momentum. We assume an efficiency for identifying electrons and muons of 85%. The events are also required to have at least one tagged b quark jet with $p_t > 30$ GeV and $|\eta| < 2.0$. The efficiency for tagging the b -jets through secondary vertices is discussed in Ref. 47, and in Chapter 4. The non- $t\bar{t}$ background coming from $Wb\bar{b}$, $Wc\bar{c}$, $Wc\bar{b}$, and $Wb\bar{c}$ final states and satisfying these criteria is small even before the H^\pm signal criteria are implemented.

Search for $H^+ \rightarrow \tau\nu$

In Method 1, we search for ℓ - τ events (e.g., $t \rightarrow bW^+ \rightarrow b\ell^+\nu$, $\bar{t} \rightarrow [\bar{b}H^- \text{ or } \bar{b}W^-] \rightarrow \bar{b}\tau^-\nu$) in which the τ decays to a single π^\pm (or K^\pm) with $p_t > 50$ GeV or with $p_t > 100$ GeV. These 1-prong decays are the most easily identified decay modes of the τ ($\tau^+ \rightarrow \pi^+\nu$ or $\tau^+ \rightarrow K^+\nu$), for which the signature is an isolated charged hadron whose momentum (from tracking) and energy (from calorimetry) should agree within errors. We selected isolated hadrons by requiring that the energy (excluding that of the pion) within a cone of radius $R = 0.4$ about the pion be less than 25% of the pion momentum. The probability for a QCD jet with $p_t > 50$ GeV to satisfy this requirement is less than 0.1% [43]. For one-prong τ candidates, the distribution of the ratio of pion track momentum to the total energy within a cone of 0.4, before making the isolation requirement, is shown in Fig. 3-54.

This study [46] was performed using the Isajet 6.31 Monte-Carlo program, including a modification to produce the correct τ polarization in the decays of W^\pm and H^\pm , as described below. The results were checked with other programs. The detector response was simulated by smearing the produced energy and momenta with the baseline calorimeter and tracking resolutions as defined in Section 3.1.1. Multiple event pile-up was not simulated, since it is not expected to affect the results at the nominal SSC luminosity. The main effect of multiple events would be to slightly decrease the efficiencies for observing isolated leptons and τ 's, but this should be a small effect for the p_t thresholds used in this study.

Implementation of the correct polarization correlations for the t quark decay chain is crucial in obtaining an accurate result. The coupling of W to $\tau\nu$ conserves chirality and the ν are left-handed, therefore (up to corrections of order M_τ/M_W) the τ^- (τ^+) are left- (right-) handed. In contrast, the τ^\pm from the decay of H^\pm would have the opposite polarizations, since the H^\pm is a scalar, and its couplings maximally violate chirality. In W decays, the τ polarization results in the preferred direction for emission of the π^\pm being opposite the momentum of the τ . In H^\pm decays, since the τ^\pm has the opposite polarization, the π^\pm tends to be emitted parallel to the τ momentum. Consequently, the isolated pion p_t spectrum from the charged Higgs decay is shifted to higher p_t . Furthermore, since we usually consider a higher mass for the H^\pm than the W , the p_t of τ 's from the H^\pm decay is already larger on average than that from W decay. In summary, both the polarization correlations and kinematic effects increase the effectiveness of the high- p_t cut on the isolated pion in enhancing the relative number of events containing a H^\pm .

The most sensitive means of detecting the presence of the charged Higgs boson decaying to τ 's is to employ lepton universality in W decays. If t quarks can only decay to W^+b , then the observed number of $\ell^+-\ell^-$ events allows us to compute the number of ℓ - τ events expected in the absence of the decay $t \rightarrow H^+b$ ($\equiv N_{\ell\tau}^{WW}$ where WW signifies that both t and \bar{t} decayed to W 's). The presence of a significant coupling of the t quark to the H^+b state would result in an excess of ℓ - τ events, independent of the theoretical calculation of the $t\bar{t}$ cross section.

If there is no charged Higgs boson, the number of $\tau^+ \rightarrow \pi^+\nu$ (or $K^+\nu$) events resulting from $t\bar{t} \rightarrow WWb\bar{b}$ with one W decaying to $\ell\nu$ (the trigger) [48] and the other to $\tau\nu$, after imposing the above-mentioned cuts and triggers, is

$$N_{\ell\tau}^{WW} = 2N_{t\bar{t}} \text{BR}(W \rightarrow \ell\nu) \text{BR}(W \rightarrow \tau\nu) e_{\ell\text{-trig}} e_{b\text{-tag}} \text{BR}(\tau \rightarrow \pi\nu) e_{\tau\pi}.$$

The variable $N_{t\bar{t}}$ is the number of $t\bar{t}$ events produced per SSC year and could be extracted from the measured value of $N_{\ell\tau}^{WW}$ if there is no H^\pm . We assume the branching ratios $\text{BR}(W \rightarrow \ell\nu)/2 = \text{BR}(W \rightarrow \tau\nu) = 1/9$ and $\text{BR}(\tau \rightarrow \pi\nu) = 11.5\%$. The variable $e_{\ell\text{-trig}}$ is the efficiency for triggering on the electron or muon (including the p_t cut, η cut, and isolation criterion), $e_{b\text{-tag}}$ is the efficiency for tagging one or both b -jets, and $e_{\tau\pi}$ is the efficiency for observing the pion (or kaon) from the τ above a minimum p_t threshold. We estimate these efficiency factors using the ISAJET Monte Carlo and our detector simulation, and list them in Table 3-11. Fig. 3-55 shows the efficiency $e_{\tau\pi}$ as a function of the p_t cut on the isolated π .

If the t quark also decays to H^+b , then we have additional final states containing $H^\pm W^\mp$ and H^+H^- , where the mixture depends on the branching ratio $B_H \equiv \text{BR}(t \rightarrow H^+b)$. The contribution from $H^+H^-b\bar{b}$ final states can be ignored, as the rate into this final state is large only when the branching ratio $\text{BR}(H^+ \rightarrow \tau\nu)$ is very small. The number of observed ℓ - τ events (from $W^\pm W^\mp b\bar{b}$ and $W^\pm H^\mp b\bar{b}$ final states) would therefore be

$$N_{\ell\tau} = N_{\ell\tau}^{WW} + N_{\ell\tau}^{WH} \equiv [(1 - B_H)^2 N_{\ell\tau}^{WW}] + [B_H(1 - B_H) N_{\ell\tau}^{WH}],$$

where

$$N_{\ell\tau}^{WH} = 2N_{t\bar{t}} \text{BR}(W \rightarrow \ell\nu) e_{\ell\text{-trig}} e_{b\text{-tag}} \text{BR}(H^+ \rightarrow \tau\nu) \text{BR}(\tau \rightarrow \pi\nu) e_{\tau\pi}.$$

For the WH decays, the estimated efficiencies are also given in Table 3-11. The efficiency for tagging the b -jets in the WH case is nearly identical to that for the WW case despite the mass difference between the H^\pm and W^\pm . As discussed below, the efficiency $e_{\tau\pi}$ for observing the pion or kaon above the p_t cut is substantially higher for the WH case than for the WW case, for any given M_{top} (see Fig. 3-55).

The parameter $\tan\beta$ of the two-doublet Higgs model enters the above equations via both B_H and $\text{BR}(H^+ \rightarrow \tau\nu)$. For illustration, in Fig. 3-56 we compare the predicted (from universality) and observed p_t spectrum of the isolated pion coming from τ decay for the case $\tan\beta \approx 1.2$ (which yields $\text{BR}(H^+ \rightarrow \tau\nu) \sim 0.50$ and $B_H \sim 0.075$) where the influence of the H^\pm would be large. In this favorable case, the excess due to the $t \rightarrow H^+b$ decays over the universality prediction is more than a factor of four.

Table 3-11

The efficiencies for lepton triggering, for b tagging (including the $p_t(b)$ cut), and for finding a π from the decay of the τ for the given $p_t(\pi)$ cut (units are GeV). The quoted efficiencies do not include the branching ratios for the t quark decays. For $M_{top} = 100, 150$, and 200 GeV, the numbers of events before efficiencies are 4×10^8 , 1×10^8 , and 4×10^7 , respectively.

M_{top} (GeV)	M_{H^\pm} (GeV)	$e_{\ell\text{-trig}}$	$e_{b\text{-tag}}$	$e_{\tau\pi}$ $p_t(\pi) > 50$	$e_{\tau\pi}$ $p_t(\pi) > 100$
100	no H^\pm	0.35	0.10	0.044	0.0022
100	75	0.34	0.13	0.16	0.020
100	85	0.35	0.087	0.18	0.04
100	95	0.36	0.066	0.27	0.077
150	no H^\pm	0.39	0.27	0.065	0.011
150	75	0.38	0.28	0.18	0.043
150	125	0.39	0.22	0.29	0.09
150	140	0.40	0.20	0.32	0.12
200	no H^\pm	0.46	0.31	0.091	0.018
200	125	0.46	0.30	0.31	0.11
200	175	0.46	0.26	0.39	0.16




FIG. 3-55. The efficiency for pions from τ decays to pass a threshold cut ($p_t(\pi) > p_0$) in WW decays (dashed) and WH decays (solid). For the lower (upper) solid histograms the charged Higgs mass is $M_{H^\pm} = 75$ GeV (125 GeV).

Since the number of $\ell^+-\ell^-$ events would be depleted by $(1 - B_H)^2$, our universality argument predicts that the number of $\ell\text{-}\tau$ events with an isolated single hadron (π^+ or K^+) is just $N'_{\ell\tau}{}^{WW}$ and the observed excess is $N_{\ell\tau} - N'_{\ell\tau}{}^{WW} = N'_{\ell\tau}{}^{WH}$. We can compute the statistical significance of this excess by comparing $N'_{\ell\tau}{}^{WH}$ to the universality prediction $N'_{\ell\tau}{}^{WW}$. In order to quantify the observability of this charged Higgs signal as a function of the Higgs model parameter $\tan\beta$, we compute the number of standard deviations by which the observed number of isolated pions exceeds the prediction from universality:

$$N_{SD} = \frac{N'_{\ell\tau}{}^{WH}}{\sqrt{N'_{\ell\tau}{}^{WH} + N'_{\ell\tau}{}^{WW}}}.$$

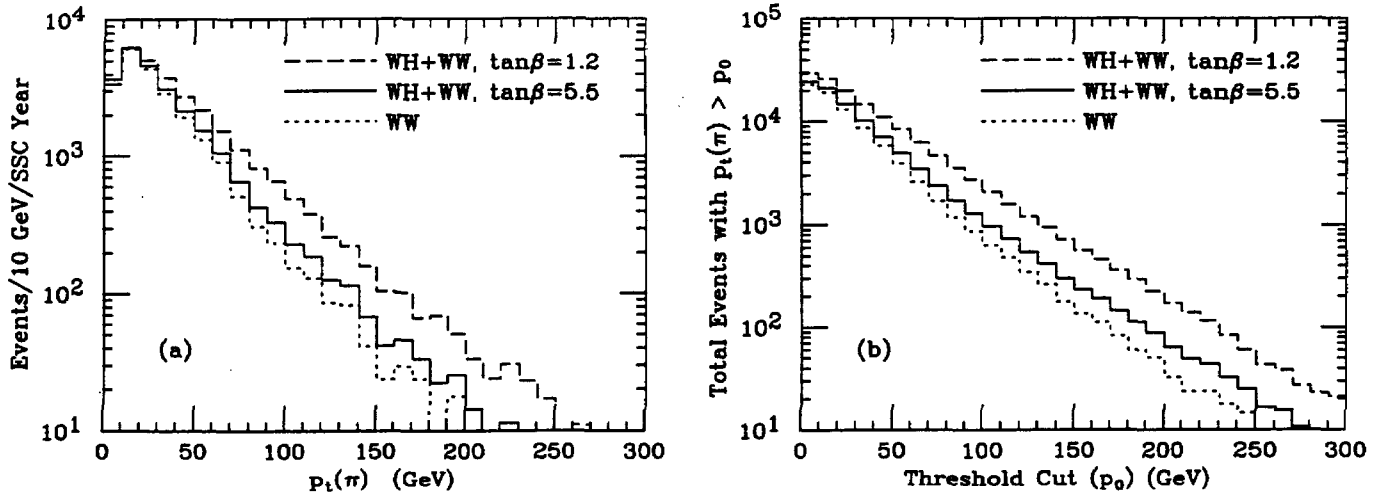


FIG. 3-56. Transverse momentum distributions for pions from $t\bar{t}$ events. (a) The transverse momentum distributions for isolated pions coming from $t\bar{t}$ events for $M_{top} = 150$ GeV. (b) The total integrated number of isolated pions with $p_t(\pi) > p_0$. In both cases the dotted histograms are the predictions from universality (i.e. corresponding to the term $N_{t\tau}^{WW}$ in the text). The solid and dashed histograms would be the actual observed spectra if a charged Higgs of mass $M_{H^\pm} = 125$ GeV is present (corresponding to the sum of $N_{t\tau}^{WW}$ and $N_{t\tau}^{WH}$). These figures employ the branching ratios predicted for $\tan\beta = 1.2$ and 5.5 .

In Fig. 3-57 and Fig. 3-58, N_{SD} is plotted for p_t cuts on the isolated pion of 50 and 100 GeV. It is critical to keep track of the polarization of the τ 's in the Monte Carlo analysis—ignoring the polarization reduces the statistical significance by a factor of two. We have estimated that the backgrounds from $Wb\bar{b}$ reduce the number of standard deviations by less than 3%. All other W -jet-jet backgrounds together are smaller than this because the b -jet tagging requirement more than compensates for the larger production cross sections of some of these channels. Requiring five standard deviations above background, we conclude that after one year of SSC running we could detect the presence in t quark decays of the charged Higgs boson decaying to τ 's for all $\tan\beta > 0.5$ as long as M_{H^\pm} is not too close to M_{top} (no less than 5–10 GeV), see Fig. 3-59. The region in parameter space with $\tan\beta \approx 5.5$ is the most marginal, since $B(t \rightarrow H^+b)$ reaches a minimum. As long as $B(t \rightarrow H^+b) \gtrsim 0.003$, the $H^+ \rightarrow \tau\nu \rightarrow \pi\nu\bar{\nu}$ decay mode can usually be used to detect the charged Higgs. For smaller values of $\tan\beta$, where $B(H^+ \rightarrow \tau\nu)$ becomes small, we must employ the $H^+ \rightarrow c\bar{s}$ decay mode, discussed below.

Although charged Higgs bosons can be easily detected via universality violation when $\tan\beta > 0.5$, the $H^\pm \rightarrow \tau\nu \rightarrow \pi\nu\bar{\nu}$ decay mode clearly makes determination of the H^\pm mass difficult. In order to measure the H^+ mass, it is necessary to estimate the momentum of the t quark and the b quark in the decay $t \rightarrow H^+b$. The two methods of Ref. 46 both rely on selecting a kinematic region where the contribution of the missing neutrinos to the H^+ mass is minimized. One method uses the momentum of the b quark and the high p_t pion to estimate the t quark momentum. The other method uses the decay products from the opposite t quark as well. Both variables seem to give some sensitivity to the charged Higgs mass, but further study is needed before reaching any conclusions.

Finally, we emphasize that the previous discussion has concentrated on a particular two-Higgs-doublet model. For comparison, we note that in other models[45] the $\tau\nu$ branching ratio is independent of $\tan\beta$, while $B(t \rightarrow H^+b)$ falls rapidly in the $\tan\beta \geq 1$ region. Thus, the use of the universality violation signal will not be possible at large $\tan\beta$. A more detailed computation of N_{SD} in this model (using the cuts and efficiencies employed above) shows that detection of the H^\pm in this model is possible for $\tan\beta \approx 0.1$ up to $\tan\beta = 2-10$, depending upon the choice of M_{H^\pm} and M_{top} .

For either model it is clear that for small values of $\tan\beta$ ($\lesssim 0.5$), where $B(H^+ \rightarrow \tau\nu)$ becomes small, we must employ the $H^+ \rightarrow c\bar{s}$ decay mode, to which we now turn.




FIG. 3-57. For three values of M_{top} the statistical significance, N_{SD} of the excess of isolated pions due to $t \rightarrow H^+b$, $H^+ \rightarrow \tau\nu$ and $\tau \rightarrow \pi\nu$ relative to expectations for $t \rightarrow W^+b$ (assuming lepton universality) as a function of $\tan\beta$. We require an isolated lepton with $p_t > 40$ GeV and an isolated pion with $p_t > 100$ GeV.




FIG. 3-58. For $M_{top} = 150$ GeV the statistical significance, N_{SD} of the excess of isolated pions due to $t \rightarrow H^+b$, $H^+ \rightarrow \tau\nu$, and $\tau \rightarrow \pi\nu$ relative to expectations for $t \rightarrow W^+b$ assuming lepton universality as a function of $\tan\beta$. We require an isolated lepton with $p_t > 40$ GeV and compare results for p_t cuts on the isolated pion of 50 GeV (solid curves) and 100 GeV (dashed curves).




FIG. 3-59. For $M_{top} = 150$ GeV and $\tan\beta = 5.5$ the statistical significance, N_{SD} of the excess of isolated pions due to $t \rightarrow H^+b$, $H^+ \rightarrow \tau\nu$, and $\tau \rightarrow \pi\nu$ relative to expectations for $t \rightarrow W^+b$ (assuming lepton universality) as a function of M_{H^\pm} . We require an isolated lepton with $p_t > 40$ GeV and compare results for p_t cuts on the isolated pion of 50 GeV (solid curve) and 100 GeV (dashed curve).

Search for $H^+ \rightarrow c\bar{s}$

In Section 3.3.2, we described a technique to determine the t -quark mass by reconstructing the hadronic decays of the t quark $t \rightarrow bW$, $W \rightarrow u\bar{d}$ (or $c\bar{s}$) in the context of Standard Model t quark decays. We have extended this technique to study a 150 GeV t quark decaying to H^+ (or W^+) with $H^+(W^+) \rightarrow u\bar{d}$ or $c\bar{s}$, with an assumed H^+ mass of 125 GeV. The ISAJET Monte Carlo was used to generate two samples of $t\bar{t}$ events with decays via W^-H^+ and W^-W^+ , which were simulated with the detector model described in Section 3.1.1. Jets were reconstructed using the clustering algorithm described in Section 3.1.1 with a cone size $R = 0.4$ in the region $|\eta| < 2.5$.

Recall that this method selects high transverse momentum events in order to simplify the combinatorics of reconstructing the correct set of three jets. In these events, one of the t quarks decays via $t \rightarrow Wb$ with $W \rightarrow \ell\nu$, and the other decays hadronically into three jets. Event rates and efficiencies for the cuts described below are listed in Table 3-12. We require an isolated electron or muon with $p_t > 40$ GeV and $|\eta| < 2.5$ so that the candidate events satisfy the inclusive lepton trigger. The lepton efficiency is identical for W^-H^+ and W^-W^+ events. We also require at least three jet clusters, each with $p_t > 30$ GeV (uncorrected) in the opposite hemisphere to the lepton ($\Delta\phi > 90^\circ$). We demand that one of these clusters have $|\eta| < 2.0$ and be tagged as a b -jet with a secondary vertex (see the discussion in Chapter 4).

As shown in Fig. 3-60, the p_t spectrum of the b -jet is much softer for the charged Higgs decay due to the small mass difference assumed between the H^+ and the t , resulting in a lower efficiency to observe the b -jet. Also, in the WH events, the two jets from the H decay are typically further apart than in WW events (Fig. 3-61).

Table 3-12
Summary of event samples and efficiencies for cuts in WW and WH events.

	WW events	WH events
$\sigma(t\bar{t})$	12 nb	
$N_{t\bar{t}}$	1.2×10^8	
Branching ratio	8/27	$4/9\text{BR}(H^+ \rightarrow c\bar{s})$
Lepton, geometric	0.43	0.43
Lepton id, isolation	0.85	0.85
b -jet, geometric	0.51	0.30
b -jet tag	0.30	0.30
$N(p_t(3\text{-jet})) > 200$ GeV	1.58×10^5	$1.26 \times 10^5\text{BR}(H^+ \rightarrow c\bar{s})$

As in the case of the Standard Model t quark search, we can reduce combinatoric backgrounds by requiring the transverse momentum of the three-jet system to be > 200 GeV. These cuts result in approximately 160,000 (125,000 $B(H^+ \rightarrow c\bar{s})$) t candidate three-jet combinations per nominal SSC year, assuming only WW (WH) decays.




FIG. 3-60. The transverse momentum distribution for b -jets opposite the lepton in $t\bar{t} \rightarrow WWb\bar{b}$ events (solid) and $t\bar{t} \rightarrow WHb\bar{b}$ events (dotted). The t quark mass is 150 GeV, and the charged Higgs mass is 125 GeV.

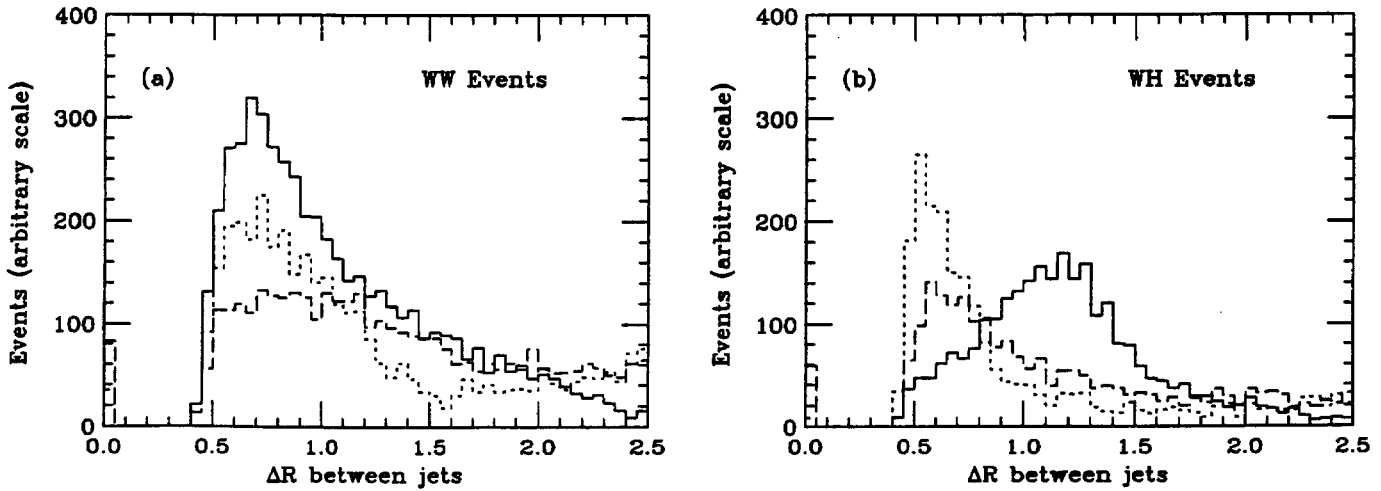


FIG. 3-61. The distance in $\eta - \phi$ space between the three jets (a) in WW events and (b) in WH events. The solid histogram is the distance between the two non- b -jets, and the dashed (dotted) histogram is the distance between the lower- p_t jet (higher- p_t jet) and the b -jet. The t quark mass is 150 GeV, and the charged Higgs mass is 125 GeV.

The invariant mass distributions for the two jets (not including the b -jet) in WW events and WH events are plotted in Fig. 3-47 and Fig. 3-62, respectively, and the three-jet invariant mass distributions are shown in Fig. 3-48 and Fig. 3-63 for the WW and WH cases. In both cases, a large invariant mass peak for the W , H , or t is evident and the combinatoric background is relatively small. As for the standard t quark decays, the peaks appear at lower invariant mass (111.4 GeV for the H) than the actual particle masses (125 GeV for the H).




FIG. 3-62. The two-jet invariant mass distribution for a sample of $t\bar{t} \rightarrow WHb\bar{b}$ events. Neither jet is the tagged b -jet. The t quark mass is 150 GeV, and the charged Higgs mass is 125 GeV. This figure should be compared with Fig. 3-47.




FIG. 3-63. The three-jet invariant mass distribution for a sample of $t\bar{t} \rightarrow WHb\bar{b}$ events. The t quark mass is 150 GeV, and the charged Higgs mass is 125 GeV. This figure should be compared with Fig. 3-48.

In subsequent plots, an average correction factor has been applied to the jet energies as a function of measured jet p_t (the dotted curve in Fig. 3-49). To further reduce the background to the H^+ signal in the two-jet invariant mass distribution, we require that the corrected three-jet invariant mass be in the range $135 < M(3\text{-jet}) < 165$ GeV. The two-jet mass distributions after this cut are shown in Fig. 3-50 and Fig. 3-64. With the jet energies corrected, the H mass peak appears at 123.3 GeV with a width of 8.5 GeV.

If we require the two-jet invariant mass in the WH case to be in the range 110 to 140 GeV, the three-jet invariant mass in Fig. 3-65 shows almost no background. The measured t invariant masses in the WW and WH modes are 147.9 GeV and 148.2 GeV, each with a width of approximately 9 GeV.




FIG. 3-64. The two-jet invariant mass distribution for samples of $t\bar{t} \rightarrow WHb\bar{b}$ events. In this plot the three-jet invariant mass is required to be consistent with the t mass by selecting events with $120 < M(3\text{-jet}) < 150$ GeV (uncorrected). The t quark mass is 150 GeV, and the charged Higgs mass is 125 GeV. This figure should be compared with Fig. 3-50.




FIG. 3-65. The three-jet invariant mass distribution for events with a W or a H^+ candidate. In this plot the two-jet invariant mass is required to be consistent with either the W or H^+ mass. The t quark mass is 150 GeV, and the charged Higgs mass is 125 GeV. This figure should be compared with Fig. 3-51.

To determine the statistical significance of the H^\pm and W^\pm mass peaks for a particular branching fraction $\text{BR}(t \rightarrow H^+b)$, the ideal technique would be to fit the dijet invariant mass distribution obtained from the data to the distributions obtained using a Monte Carlo that produces both WW and WH events (with a W or H decaying to jets). This will give an estimate for the fraction of WH events in the data (or a limit on the amount of WH present). This determines the product of $\text{BR}(t \rightarrow bH^+)$ and $\text{BR}(H^+ \rightarrow c\bar{s})$, so that $\tan\beta$ can be determined using Fig. 3-52 and Fig. 3-53. As an example, the two-jet invariant mass distributions obtained after imposing the selection criteria (and efficiencies) are shown in Fig. 3-66 for the particular cases of $\tan\beta = 0.4$ (yielding $B_H = 0.42$ and $\text{BR}(H^+ \rightarrow c\bar{s}) = 0.98$) and $\tan\beta = 1.0$ (yielding $B_H = 0.096$ and $\text{BR}(H^+ \rightarrow c\bar{s}) = 0.59$). For $\tan\beta = 0.4$, the W^\pm and H^\pm mass peaks are both very prominent, and discovery of the H^\pm is clearly possible. The $\tan\beta = 1.0$ case is somewhat marginal, since the combinatoric background would have to be well understood to claim a signal. The statistical significance is quite high since the statistical errors are small; however, in this case we would prefer to rely on Method 1 ($H^+ \rightarrow \tau\nu$) which is effective down to $\tan\beta = 0.5$.

Since our mass assumptions yield two reasonably well separated mass peaks, we have adopted a simple technique to estimate the statistical significance of each peak. We consider the number of events in the two intervals $65 < M_{jj} < 95$ GeV (W interval) and $110 < M_{jj} < 140$ GeV (H interval). The background beneath the H peak has a contribution from both WW and WH events. Unlike the case in the previous study [25], the shapes of the background beneath the peaks in the WW and WH events are dissimilar. We have estimated the level of background under the H peak in WH events by using a smooth curve overlapping each histogram on both edges of the H interval, and we expect this estimate to have an uncertainty of $\sim 10\%$. We follow the converse procedure in determining the number of events in the W^\pm peak. For the two mass intervals, the estimated ratios of signal and of background events to the number of events passing the trigger and event selection requirements in the WW and WH processes are shown in Table 3-13.

Table 3-13

The fraction of events, relative to the total number of events passing the trigger and event selection requirements, in different mass regions. These regions are defined to be $65 < M(2\text{-jet}) < 95$ GeV (for WW events) or $110 < M(2\text{-jet}) < 140$ GeV (for WH events). All events were required to satisfy $135 < M(3\text{-jet}) < 165$ GeV. The background beneath a given peak has contributions from both WW and WH events.

	WW Events	WH Events
W peak signal	0.312	0
H peak signal	0	0.296
W peak background	0.103	0.059
H peak background	0.0185	0.055

To quantify the statistical significance of the H^\pm and W^\pm mass peaks, we plot the number of standard deviations above background as a function of $\tan\beta$:

$$N_{SD} = \frac{N_{\text{above}}}{\sqrt{N_{\text{above}} + N_{\text{below}}}},$$

where N_{above} is the number of "excess" events appearing above the background curve and N_{below} is the number below the background curve in the two mass intervals mentioned above. The resulting values for N_{SD} are plotted as a function of $\tan\beta$ in Fig. 3-67. The highest $\tan\beta$ value for which we could discover the charged Higgs by this method would depend critically on understanding the shape of the combinatoric background. To be conservative we should claim to see a signal only when the shape of the distribution is clearly different from the background. Hence, we argue that this method of H^\pm detection is valid only for $\tan\beta < 1.0$.




FIG. 3-66. The two-jet invariant mass distribution after two-jet selection criteria are imposed, for the cases of (a) $\tan\beta = 0.4$ and (b) $\tan\beta = 1.0$. The plot is normalized to one SSC year of running. The dotted curve in (b) indicates the background determined as described in the text. The figures have been numerically smoothed to more accurately convey the statistical power of the data. The error bars shown correspond to the statistics for one SSC year. The t quark mass is 150 GeV and the H mass is 125 GeV.




FIG. 3-67. The statistical significance, N_{SD} , of the charged Higgs and W boson (dashed) peaks in the two-non- b -jet invariant mass distribution as a function of $\tan\beta$. Also shown is the statistical significance obtained using the $\tau\nu$ channel, where the solid (dot-dashed) curves are for $p_t(\pi) > 50$ (100) GeV. We assume one SSC year of running and have taken $M_{top} = 150$ GeV and $M_{H^\pm} = 125$ GeV.

Summary of charged Higgs search

We have examined charged Higgs boson production in $t\bar{t}$ events in which one t decays to $H^\pm b$ and the other to $W^\pm b$. This study was performed in the context of a two-Higgs doublet model in which one Higgs doublet couples only to up-type quarks and the other only to down-type quarks and to leptons. In the particular case of $M_{top} = 150$ GeV and $M_{H^\pm} = 125$ GeV, discovery of the charged Higgs boson will be possible over the entire interesting range of parameter space. The detection of H^\pm for $\tan\beta \gtrsim 0.5$ is possible through $H^+ \rightarrow \tau\nu$ decays, while for $\tan\beta \lesssim 1.0$ the decay $H^+ \rightarrow c\bar{s}$ can be employed. This study illustrates the importance of efficiently identifying b -quark jets and τ 's.

3.4. SUSY searches

3.4.1. Supersymmetric phenomenology

Supersymmetry offers the possibility of a consistent unification of particle physics and gravity. In theories of "low-energy" supersymmetry, the effective scale of supersymmetry breaking is tied to the electroweak scale[49]. In this way, supersymmetry may ultimately explain the origin of the large difference between the W and Z masses and the Planck scale.

The minimal supersymmetric extension of the Standard Model (MSSM) consists of taking the Standard Model and adding the corresponding supersymmetric partners [50]. In addition, the MSSM contains two Higgs doublets, which have been discussed in Section 3.2.4 and Section 3.3.3. Supersymmetric interactions consistent with global $B - L$ conservation (baryon minus lepton number) are assumed. Finally, the most general soft-supersymmetry-breaking terms are added. If supersymmetry is relevant for explaining the scale of electroweak interactions, then the mass parameters that occur in the soft-supersymmetry-breaking terms must be of order 1 TeV or less.

As a consequence of $B - L$ invariance, the MSSM possesses a discrete R -parity invariance, where $R = (-1)^{3(B-L)+2S}$ for a particle of spin S [53]. This implies that all the ordinary Standard Model particles have even R -parity, whereas the corresponding supersymmetric partners have odd R -parity. The conservation of R -parity in scattering and decay processes has a crucial impact on supersymmetric phenomenology. For example, starting from an initial state involving ordinary (R -even) particles, it follows that supersymmetric particles must be produced in pairs. In general, these particles are highly unstable and decay quickly into lighter states. However, R -parity invariance also implies that the lightest supersymmetric

particle (LSP) is absolutely stable, and must eventually be produced at the end of a decay chain of a heavy unstable supersymmetric particle. In order to be consistent with cosmological constraints, the LSP is almost certainly electrically and color neutral[54]. Consequently, the LSP is weakly interacting in ordinary matter, i.e. it behaves like a neutrino and will escape detectors without being directly observed. Thus, the canonical signature for (transverse) energy, due to the escape of the pair of LSP's. Some model builders attempt to relax the assumption of R -parity conservation[55]. Models of this type must break $B - L$ and are therefore strongly constrained, though not presently ruled out.

The parameters of the MSSM fall into two classes: a supersymmetry-conserving sector and a supersymmetry-breaking sector. Among the parameters of the supersymmetry-conserving sector is a Higgs mass parameter μ . The supersymmetry-violating sector contains various parameters including (i) gaugino Majorana masses M_3 , M_2 and M_1 associated with the SU(3), SU(2) and U(1) subgroups of the Standard Model and (ii) three scalar Higgs mass parameters that can be re-expressed in terms of the two Higgs vacuum expectation values, v_1 and v_2 , and one physical Higgs mass. Here, v_1 (v_2) is the vacuum expectation value of the Higgs field that couples exclusively to down-type (up-type) quarks and leptons. Note that $v_1^2 + v_2^2 = (246 \text{ GeV})^2$ is fixed by the W mass (or equivalently by the Fermi constant G_F), while the ratio

$$\tan \beta = v_2/v_1 \quad (3.1)$$

is a free parameter of the model.

The supersymmetric partners of the gauge and Higgs bosons are fermions. The *gluino* is the color octet Majorana fermion partner of the gluon with mass $M_{\tilde{g}} = |M_3|$. The supersymmetric partners of the electroweak gauge and Higgs bosons (the *gauginos* and *Higgsinos*) can mix. As a result, the physical mass eigenstates are model-dependent linear combinations of these states, called *charginos* and *neutralinos*, which are obtained by diagonalizing the corresponding mass matrices. The chargino mass matrix depends on M_2 , μ , $\tan \beta$ and M_W [56]. The corresponding chargino mass eigenstates are denoted by $\tilde{\chi}_1^+$ and $\tilde{\chi}_2^+$, where the states are ordered such that $M_{\tilde{\chi}_1^+} \leq M_{\tilde{\chi}_2^+}$. The neutralino mass matrix depends on M_1 , M_2 , μ , $\tan \beta$, M_Z and the weak mixing angle θ_W [56]. The corresponding neutralino eigenstates are usually denoted by $\tilde{\chi}_i^0$ ($i = 1, \dots, 4$), according to the convention that $M_{\tilde{\chi}_1^0} \leq M_{\tilde{\chi}_2^0} \leq M_{\tilde{\chi}_3^0} \leq M_{\tilde{\chi}_4^0}$. If M_1 and M_2 are small compared to M_Z (and μ), then the lightest neutralino $\tilde{\chi}_1^0$ will be nearly a pure photino, $\tilde{\gamma}$ (the supersymmetric partner of the photon). It is common practice in the literature to reduce the supersymmetric parameter freedom by requiring that all three gaugino mass parameters are equal at some grand unification scale. Then, at the electroweak scale, the gaugino mass parameters can be expressed in terms of one of them. Having made this assumption, the chargino and neutralino masses and mixing angles depend only on three unknown parameters: the gluino mass, μ , and $\tan \beta$.

The supersymmetric partners of the quarks and leptons are spin-zero bosons: the *squarks*, charged *sleptons* and *sneutrinos*. For a given fermion f , there are two supersymmetric partners \tilde{f}_L and \tilde{f}_R , which are scalar partners of the corresponding left and right-handed fermion. (There is no $\tilde{\nu}_R$.) However, in general, \tilde{f}_L and \tilde{f}_R are not mass eigenstates since there is \tilde{f}_L - \tilde{f}_R mixing. Due to the appearance of the *fermion* mass in the corresponding element of the scalar mass-squared-matrix, one expects M_{LR} to be small compared to the diagonal squark and slepton masses, with the possible exception of the t -squark, since M_{top} is large.

In "low-energy" supergravity or "superstring-inspired" models, five flavors of squarks (with two squark eigenstates per flavor) are nearly mass-degenerate and somewhat heavier than six flavors of degenerate sleptons (with two per flavor for the charged sleptons and one per flavor for the sneutrinos). The t -squark masses are sensitive to the strength of the \tilde{t}_L - \tilde{t}_R mixing.

In this section we will discuss the signals for gluinos (\tilde{g}), squarks (\tilde{q}), charginos ($\tilde{\chi}_i^\pm$), and neutralinos ($\tilde{\chi}_i^0$). The phenomenology of the neutral Higgs bosons (H_1^0 and H_2^0) is discussed in Section 3.2, and that of the charged Higgs bosons (H^\pm) is discussed in Section 3.3.

Gluinos and squarks

The discovery of gluinos and squarks and the exploration of their properties will require the study of several signals. Gluinos and squarks can be produced in pairs ($\tilde{g}\tilde{g}$ or $\tilde{q}\tilde{q}$) or in combination ($\tilde{g}\tilde{q}$). It is not known whether the gluino or squark is lighter. The heavier one will be much harder to discover, because it will simply decay into the lighter one ($\tilde{g} \rightarrow q\tilde{q}$ or $\tilde{q} \rightarrow q\tilde{g}$). Therefore the signals will be similar, but will contain additional jets. Such jets can also come from initial- or final-state radiation. Whether or not the heavier particle can be isolated may depend on the size of the mass splitting.

Let us begin with the gluino, assuming that $M_{\tilde{g}} < M_{\tilde{q}}$. There are many possible signatures for gluinos [57] depending on $M_{\tilde{g}}$ and the masses of the charginos and neutralinos. Taking for example $M_{\tilde{g}} = 750$ GeV and the parameters $\mu = 150$ GeV and $\tan\beta = 1.5$, we find branching ratios of

$\tilde{g} \rightarrow \tilde{\chi}_1^\pm \bar{q}q'$	30%
$\tilde{g} \rightarrow \tilde{\chi}_2^\pm \bar{q}q'$	30%
$\tilde{g} \rightarrow \tilde{\chi}_1^0 \bar{q}q$	10%
$\tilde{g} \rightarrow \tilde{\chi}_2^0 \bar{q}q$	17%
$\tilde{g} \rightarrow \tilde{\chi}_4^0 \bar{q}q$	13%

The charginos in this case can decay via:

$\tilde{\chi}_1^\pm \rightarrow \tilde{\chi}_1^0 \bar{q}q'$	67%
$\tilde{\chi}_1^\pm \rightarrow \tilde{\chi}_1^0 \ell\nu$	33%

and

$\tilde{\chi}_2^\pm \rightarrow \tilde{\chi}_1^\pm Z^0$	40%
$\tilde{\chi}_2^\pm \rightarrow \tilde{\chi}_1^\pm H_1^0$	5%
$\tilde{\chi}_2^\pm \rightarrow \tilde{\chi}_1^0 W^\pm$	10%
$\tilde{\chi}_2^\pm \rightarrow \tilde{\chi}_2^0 W^\pm$	30%
$\tilde{\chi}_2^\pm \rightarrow \tilde{\chi}_3^0 W^\pm$	15%

and similarly for the four neutralinos.

Obviously many signatures are possible. Even for $M_{\tilde{g}} = 750$ GeV the $\tilde{g}\tilde{g}$ cross section is large, yielding 7×10^5 events per year, so that even small branching ratios may be visible. We assume that $\tilde{\chi}_1^0$ is the LSP and is unobservable in the detector; when it is produced directly in the decay of the gluino (rather than in a sequential decay), large missing transverse energy can occur. The production of $\tilde{g}\tilde{g}$ can result in events containing:

- Large missing transverse energy and jets,
- Leptons and jets,
- Leptons, jets, and large missing transverse energy,
- Z bosons and jets,
- Z bosons, jets, and large missing transverse energy.

While large missing transverse energy is a vital indicator for supersymmetry, confirmation that this missing- E_T signal is due to gluinos will require the use of additional signatures. The best of these complementary signatures consists of events with same-sign dileptons that also have energetic jets. As we show below, such events are background-free and therefore ideal for determining the gluino mass. The SDC detector provides charge identification for both electrons and muons and is therefore an effective tool in searching for gluinos.

The signals for a heavy gluino will be quite dramatic in several modes. In the following discussion, we will concentrate instead on the more difficult case of a 300 GeV gluino. For much of parameter space the branching ratios of the gluino are: $B(\tilde{\chi}_1^\pm \bar{q}q') \approx 60\%$, $B(\tilde{\chi}_1^0 \bar{q}q) \approx 15\%$, and $B(\tilde{\chi}_2^0 \bar{q}q) \approx 25\%$.

3.4.2. The missing- E_t signature

Phenomenology

If a pair of gluinos is produced, the probability that both gluinos decay directly to $\tilde{\chi}_1^0$ is rather small (about 2%). Instead we consider the process in which one gluino decays as

$$\tilde{g} \rightarrow \tilde{\chi}_1^0 \bar{q} q$$

and the other decays as

$$\tilde{g} \rightarrow \tilde{\chi}_1^\pm \bar{q} q' \text{ or } \tilde{\chi}_2^0 \bar{q} q ,$$

and then $\tilde{\chi}_1^\pm$ or $\tilde{\chi}_2^0$ decays via modes such as $\tilde{\chi}_1^0 \bar{q} q'$. The resulting events from

$$\tilde{g} \tilde{g} \rightarrow q \bar{q} q \bar{q} q \tilde{\chi}_1^0 \tilde{\chi}_1^0$$

(or $\ell^+ \ell^- q \bar{q} q \tilde{\chi}_1^0 \tilde{\chi}_1^0$ or $\ell \nu q \bar{q} q \tilde{\chi}_1^0 \tilde{\chi}_1^0$) have missing transverse energy, and typically 3–6 jets depending on the energy cuts and the coalescing algorithm. As shown in Fig. 3-68, this yields missing- E_t comparable to that from the case in which both gluinos decayed directly to $\tilde{\chi}_1^0$.




FIG. 3-68. Comparison of the resulting missing- E_t when both gluinos decay directly to the LSP ($\tilde{\chi}_1^0$) and when only one decays directly (the other undergoes a cascade decay). The final state is three or more jets each with $E_t > 70$ GeV separated by $\Delta R = 0.7$. The jets are required to have $|\eta| < 5$. Events are rejected if they contain a jet with $E_t > 70$ GeV within an azimuthal angle of 20° of the missing- E_t or if they have circularity (defined in the text) $C < 0.2$. The 2-LSP curve has been normalized to have the same area as the 1-LSP curve to emphasize the similarity in shape.

Discovering a 300 GeV gluino using a missing- E_t signature requires the ability to remove detector-dependent backgrounds originating from mismeasured multijet events. It also requires methods to isolate detector-independent backgrounds due to production of $Z(\rightarrow \nu\bar{\nu}) + \text{multijets}$, and of $b\bar{b}$ and $t\bar{t}$ pairs with semileptonic decays.

To isolate the missing- E_t signature for gluinos we require at least three jets with $p_t > 70$ GeV, since gluino pair production gives four quarks in the initial decays. As discussed below, the primary detector-dependent background is from multijet events in which some jets are mismeasured due to resolution or cracks (Ref. 59) or in which a high- p_t jet is produced beyond the η coverage of the forward calorimeter. The observation of $E_t^{\text{miss}} > 100$ GeV due to mismeasurement of a jet or from an $E \gtrsim 7$ TeV jet produced at $|\eta| > 5$ is an extremely rare occurrence. But this small probability coupled to the large rate for QCD multijet events can lead to appreciable backgrounds to missing- E_t signatures. We first discuss the effects of energy mismeasurement; the background due to jets produced at large η will be discussed later in the context of the forward calorimeter requirements.

Missing- E_t from mismeasurement

In the case of mismeasurement, we find that the resulting background event rate is small compared to the gluino signal until nonGaussian tails are introduced into the jet energy resolution. We have studied the impact of substantial nonGaussian tails, using a model developed to describe the CDF detector at Fermilab [58]. There is not yet any similar parametrization of the response of the SDC calorimeter system, but it is expected to be substantially more Gaussian than that of CDF because the SDC detector is much more hermetic (all intermodule gaps are minimized, and in many cases nonprojective). Nevertheless, we have used this model of jet energy mismeasurement to explore the potential effects on the physics capabilities. The parametrization of nonGaussian tails assumed in our studies is shown in Fig. 3-69.

We find that it is very unlikely that two jets will be badly mismeasured in a single event (since large mismeasurement is rare). Furthermore, it is much more probable to measure an $E_t = 800$ GeV jet to be 700 GeV, than to measure an $E_t = 300$ GeV jet to be 200 GeV (both yield 100 GeV of missing- E_t). The distributions of the mismeasured background events look quite different from those of the signal. The missing- E_t vector will clearly lie quite close to one of the jets. Therefore, the background distribution of the azimuthal angle ϕ between the missing- E_t vector and the nearest jet with $E_t \geq 70$ GeV is peaked at zero unlike the signal (gluino pair production). The E_t of the leading and non-leading jets in this background are also much larger than those in the signal. Another effective variable is the circularity C , which is defined as $C = \frac{1}{2} \min(\sum \vec{E}_t \cdot \hat{n})^2 / (\sum E_t^2)$, where the sum is over calorimeter cells and where the minimization is over all \hat{n} (a unit vector in the transverse plane); $C = 0$ corresponds to pencil-like events while $C = 1$ corresponds to isotropic events.

In practice, it is only necessary to cut on the ϕ variable. Figure 3-70 shows the missing- E_t distribution from gluino pair production and decay, and also from the background processes. The distributions are for events with $\phi > 20^\circ$. This cut eliminates much of the background due to nonGaussian tails on the jet energy resolution.

The hadron calorimeter jet resolution can be parametrized as $\Delta E/E = A/\sqrt{E} \oplus B$ where we expect to achieve $A \approx 60\%$ and $B \approx 2\%$ for $|\eta| < 3$ (see Section 3.1.1). If we double this resolution (both A and B), there is a modest impact on the mismeasurement background. The missing- E_t processes do not severely constrain the hadron calorimeter resolution.

There are other backgrounds that do not depend on detector parameters. One such detector-independent background is Z plus multijets production, where $Z \rightarrow \nu\bar{\nu}$. This background is small. However, $b\bar{b}$ and $t\bar{t}$ production with semileptonic decays lead to substantial backgrounds. Important cuts here include the circularity and the ϕ variable described above. As shown in Fig. 3-70, we are able to isolate the gluino signal for missing- $E_t > 100$ GeV, yielding something like a million events per year with the cuts described in the figure caption. The $Z + jets$ and $t\bar{t}$ rates will be measured by the SDC detector and therefore the backgrounds can be subtracted. The background from mismeasured jets can be significantly reduced by further cuts.

In the above analysis we have included the effects of initial-state radiation (which affects signals and backgrounds similarly) and of neutrinos and missing muons (for $|\eta| > 2.5$). Because we consider only large values of missing- E_t (missing- $E_t > 100$ GeV), we expect pileup to be unimportant in this analysis; even if it were important, this search could be done at one tenth of design luminosity because of the extremely high event rate.




FIG. 3-69. The comparison of a Gaussian jet energy resolution function (dashed curve) with a resolution (solid curve) in which nonGaussian tails have been included. E_{in} is the energy of the jet before mismeasurement (fixed at 500 GeV for this figure), while E_{out} represents the measured energy. It is expected that the SDC detector will have significantly smaller tails than those shown here.




FIG. 3-70. Search for evidence of gluino pair production in the distribution of missing- E_t for the final state of three or more jets each with $E_t > 70$ GeV separated by $\Delta R = 0.7$. The jets are required to have $|\eta| < 5$. Events are rejected if they contain a jet with $E_t > 70$ GeV within an azimuthal angle of 20° of the missing- E_t vector or if they have circularity (defined in the text) $C < 0.2$. The signal is for pair production of 300 GeV gluinos decaying as described in the text. The lower solid histogram is the sum of the individual background contributions, the upper solid histogram is the sum of the background and the gluino signal. The detector-dependent background due to multijet events with missing- E_t generated by calorimeter resolution or by energy loss out of the end of the detector ($|\eta| > 5$) is shown as a dashed curve. The dotted background arises from the final states $t\bar{t}$ and $b\bar{b}$, where the missing- E_t is due to neutrinos. The dash-dotted background is due to $Z + multijet$ events.

3.4.3. Squarks and missing- E_t

The discussion of squark decays is more complicated than that for gluino decays, because there are both \tilde{q}_L and \tilde{q}_R , and because there are six flavors of squarks. Although all have similar types of decays, each of these can have very different branching ratios. If $M_{\tilde{q}} < M_{\tilde{g}}$, then the general form of the decays is $\tilde{q} \rightarrow q' + \tilde{\chi}_i$. We expect the \tilde{q}_L to decay as:

$$\tilde{q}_L \rightarrow q' \tilde{\chi}_i^\pm$$

or

$$\tilde{q}_L \rightarrow q \tilde{\chi}_i^0.$$

However, the \tilde{q}_R cannot decay into charginos just as right-handed quarks have no coupling to the W boson (note that the t -squark, \tilde{t} , is an exception). The only decay for \tilde{q}_R is

$$\tilde{q}_R \rightarrow q \tilde{\chi}_i^0.$$

The resulting charginos and neutralinos decay in the same manner as discussed for gluinos. The cascade decays of squarks again lead to a variety of signatures. For a wide range of parameters (for $|\mu| \gtrsim M$) the \tilde{q}_R decay dominantly to $\tilde{\chi}_1^0$. This happens in part because the usual dominant decay to charginos

is not allowed. By averaging over the production of \tilde{q}_L and \tilde{q}_R and over flavors, one finds [60] that the branching ratio for $\tilde{q}\tilde{q} \rightarrow q\tilde{q}\tilde{\chi}_1^0\tilde{\chi}_1^0$ is about 40% for $|\mu| \gtrsim M$. For smaller values of $|\mu|$, one can rely on the mode in which only one squark decays directly to $\tilde{\chi}_1^0$, which is still substantial, 10–20%, and leads to large missing- E_t . The signature is of course different in that two hard jets are produced instead of the four jets from $\tilde{g}\tilde{g}$ production although there is somewhat more missing- E_t .

In summary, the cross sections for squark production are somewhat smaller than those for gluino production so that the range of masses covered by the missing- E_t signature may be slightly smaller than the 300 to 1000 GeV range that can be covered in the gluino case.

3.4.4. Missing- E_t and the forward calorimeter

Without the forward calorimeter, jets would be detected only out to $|\eta| = 3$ ($\theta = 5.7^\circ$). A jet with an energy of 1 TeV at $|\eta| = 3$ would have a transverse energy of 100 GeV. If that jet were not detected, the event would appear to have missing- E_t of 100 GeV. Such events are extremely common (there are almost 10^{10} such events per year at SSC design luminosity—see Section 3.7), and would make it impossible to look for signatures involving missing- E_t for Higgs bosons, new heavy quarks, and supersymmetric particles such as gluinos, squarks, charginos, and neutralinos. One example of this has been already shown in the context of a heavy Higgs search (see Fig. 3-27).

Clearly the missing- E_t capabilities of the SDC detector discussed previously are crucial for a variety of physics processes in which one or more noninteracting particles (such as neutrinos and photinos) are produced. However, perhaps the most difficult challenge to the detector's capabilities would arise in the search for light gluinos (with say $M(\tilde{g}) = 300$ GeV). Such gluinos produce only 100–200 GeV of missing- E_t , yet they have the largest possible detector-dependent background. As discussed previously, the signature for a light gluino is the observation of events with three or more energetic jets and with missing- $E_t > 100$ GeV. The detector-dependent background is multijet QCD events in which one or more jets is mismeasured or escapes out the end of the detector. In the simplest approximation, the forward calorimeter discriminates against such backgrounds by observing jets in the forward ($|\eta| > 3$) region that would otherwise escape undetected (yielding fake missing- E_t). However, merely vetoing such events would not be adequate to reduce backgrounds below the light gluino signal. In fact, most high p_t events will contain initial-state radiation in the forward region, so that vetoing will simply lower the rate for both the signal and the background.

Before proceeding to discuss the forward calorimeter requirements in more detail, it is worthwhile to briefly consider the kinematics of the forward region. This region corresponds to very small angles from the beam: $|\eta| = 3$ is 5.7° while $|\eta| = 5$ is 0.8° . What is the size of a jet in this region? In the central region (near $\eta = 0$) we know that high p_t jets are approximately $\Delta\eta = 0.5$ and $\Delta\phi = 0.5$ in size, or equivalently $\Delta\theta = 30^\circ$ and $\Delta\phi = 30^\circ$. Two-jet production of $E_t = 100$ GeV jets in the central region has $x_1 \approx x_2$, and a parton-parton energy, $\hat{s} = s x_1 x_2$, which is small. At $|\eta| = 4$ an $E_t = 100$ GeV jet has $E = 2.7$ TeV, so that if $x_1 \approx x_2$, then $\hat{s} = s x_1 x_2 = (5.4 \text{ TeV})^2$. In this case because \hat{s} , x_1 , and x_2 are all large, the cross section for these two-jet events, which are at rest in the lab, is strongly suppressed. Instead, the dominant source for jets in this region of η is a boosted scattering (x_1 small, x_2 large or vice versa). In this case, the center-of-mass energy \hat{s} is small, and the cross section is not suppressed. Furthermore, in this case we find $\Delta\eta = 0.5$ and $\Delta\phi = 0.5$ so that jets have $\Delta\theta = 1^\circ$ and $\Delta\phi = 30^\circ$. These results are confirmed [61] by parton-shower Monte Carlo studies.

We return to the case of the light gluino and ask what coverage in η is required. If the forward calorimeter covers the region out to $|\eta| = \eta_f$, then we are looking for events with three or more jets each with $E_t > 70$ GeV and $|\eta(\text{jet})| < 3$ plus an additional jet with $|\eta(\text{jet})| > \eta_f$. This additional jet leads to fake missing- E_t . From Fig. 3-71 we see that $\eta_f = 4$ is inadequate while $\eta_f = 5$ reduces the background well below the signal for missing- $E_t > 100$ GeV. This study used a parton-level calculation, so that the result should be interpreted in terms of the fiducial coverage for the forward calorimeter—additional coverage will be required to collect the particles from a jet whose axis is at $\eta = 5$.

If the jet resolution in the forward calorimeter is parametrized as $\Delta E/E = A/\sqrt{E} \oplus B$, then the stochastic term, A , is of little relevance due to the large energy values which appear in this region. By comparing the QCD multijet background to the light gluino signal, our studies show that $B = 20\%$ yields marginal results while $B = 10\%$ is sufficient. It should be emphasized that these are the resolutions in the forward calorimeter only.

Next, we turn to the question of segmentation in the forward calorimeter. At the most naive level, the segmentation need only be adequate to measure the missing- E_t accurately; no reconstruction of jets is needed to do this. This requirement would suggest that relatively crude segmentation would be adequate (perhaps $\Delta\phi = \Delta\eta = 0.4$). However, if we do not identify and reconstruct jets, we cannot make the ϕ cut, and we find that the background to the light gluino signal is an order of magnitude larger than the signal. This is shown in Fig. 3-72. As discussed above, we expect the background to arise from significant mismeasurements of one and only one very high E_t jet. If we are able to efficiently detect this jet in the forward calorimeter, then we can eliminate all events in which there is a jet within a small azimuthal angle of the missing- E_t vector. This is a powerful cut and brings the background well under the signal. The critical point here is that the segmentation (and resolution) must be adequate to identify this jet and to accurately measure its separation from the missing- E_t vector. The conclusion of both parton Monte Carlo studies (as in Fig. 3-72) and of ISAJET studies is that segmentation of $\Delta\eta = \Delta\phi = 0.2$ is required in the forward calorimeter to assure adequate rejection of missing- E_t backgrounds (recall that the requirement on the ϕ separation between the two vectors was typically 20° to 40° in the preceding analyses; this leads to a segmentation requirement which is slightly finer than the actual cut).




FIG. 3-71. Comparison of the missing- E_t distributions for the background (to light gluino pair production) due to multijet events with energy loss out of the end of the detector, $|\eta| > 4$ (dashed histogram) or 5 (solid histogram). The events are required to have three or more jets each with $E_t > 70$ GeV and $|\eta| < 3$ separated by $\Delta R = 0.7$. Events are rejected if they contain a jet with $E_t > 70$ GeV within an azimuthal angle of 40° of the missing- E_t . The solid curve is for pair production of 300 GeV gluinos decaying as described in the text.




FIG. 3-72. Comparison of the missing- E_t distributions for the background (to light gluino pair production) due to multijet events with mismeasurement of a jet. The two histograms are with (solid) and without (dashed) a cut on events containing a jet with $E_t > 70$ GeV within an azimuthal angle of 40° of the missing- E_t . The events are required to have three or more jets each with $E_t > 70$ GeV and $|\eta| < 3$ separated by ΔR of 0.7. The solid (dotted) curve is for pair production of 300 GeV (500 GeV) gluinos decaying as described in the text.

In summary, it is essential to be able to discriminate against events in which a jet is pointing in the same direction as the missing- E_t , since this cut provides a powerful tool to eliminate background with little impact on signal. Thus, the forward calorimeter must be designed not only to measure missing- E_t with suitable precision, but also to identify jets lying near the missing- E_t vector.

In the preceding discussion, we have concentrated on a gluino of mass 300 GeV. For comparison we have also shown the distribution for a 500 GeV gluino in Fig. 3-72. In searching for such a gluino, one would consider more stringent cuts, *e.g.*, requiring three jets with $E_t > 100$ GeV and a fourth jet with $E_t > 75$ GeV. This would not have much impact on the signal but would be more effective at reducing backgrounds than the cuts made for Fig. 3-72.

3.4.5. Same-sign lepton signature for gluinos and squarks

As discussed above in Section 3.4.1, an excellent signal for gluinos [62] is the presence of events with isolated high- p_t same-sign dileptons plus at least 4 jets. This signature can be used for the discovery or confirmation of gluinos and allows a measurement of the mass of the gluino. The same-sign events arise because the gluino is a Majorana particle and can therefore decay to charginos of *either* charge: $\tilde{g} \rightarrow q\bar{q}'\tilde{\chi}_i^\pm$. The charginos in turn can decay to $\ell\nu\tilde{\chi}_1^0$. The full process is then

$$\tilde{g}\tilde{g} \rightarrow \ell^\pm\ell^\pm q\bar{q}q\bar{q}X.$$

Since typical branching ratios are: $B(\tilde{g} \rightarrow q\bar{q}'\tilde{\chi}_i^\pm) \approx 60\%$ and $B(\tilde{\chi}_i^\pm \rightarrow \ell\nu\tilde{\chi}_1^0) \approx 20\%$, the full process occurs with a branching ratio of about 2%. The rate for this process is rather independent of the SUSY parameters. Half of these final states will be same-sign ($\ell^+\ell^+X$ or $\ell^-\ell^-X$) dilepton events. Before applying the cuts given below, a 180 GeV gluino would yield roughly 2×10^6 same-sign dilepton events in an SSC year, while a 2 TeV gluino would give about 25 same-sign dilepton events per year. In the following discussion, we continue to use a gluino with a mass of 300 GeV.

We accept only events with four jets with $|\eta| < 3$ having $p_t > 100, 60, 60,$ and 50 GeV. Each lepton is required to be isolated and to have $p_t > 20$ GeV or alternatively one lepton must have $p_t > 40$ GeV and the other $p_t > 15$ GeV. The isolation requirement consists of demanding that the leptons be separated from the four highest p_t jets by $\Delta R > 0.5$. For the present study, we employ a parton-level Monte Carlo including resolution smearing but no fragmentation or initial and final-state radiation.

The primary background is from $t\bar{t}$ events in which the \bar{t} decays to $\bar{b}\ell^-\nu$ and the t decays to b , which in turn decays to $c\ell^-\nu$ (or vice versa). If we were to use opposite-sign leptons, both leptons could originate from the W in the $t \rightarrow W + X$ decay, and the isolation and p_t cuts would not be effective in eliminating the background.

Fig. 3-73 shows the combined mass of the lower p_t lepton and the two nearest jets (of the four highest p_t jets). In the $t\bar{t}$ background the lower p_t lepton is presumably the one from the b decay. This figure, which has not had an isolation cut applied to the background, already clearly separates the signal and background. An isolation cut would independently eliminate most of the $t\bar{t}$ signal, since the lepton from the b decay almost always has a nearby jet.

While the above distribution already gives a reasonably accurate determination of the gluino mass, it can be measured more precisely by finding the invariant mass of the highest p_t lepton with the two nearest jets (those two of the four highest p_t jets which have the smallest ΔR relative to the lepton). In this case, we require the highest p_t lepton to have $p_t > 65$ GeV to minimize the missing- E_t . From the clear separation between the invariant mass peaks for the two cases shown in Fig. 3-74, we see that this procedure allows the determination of the gluino mass with a precision of $\pm 10\%$.

The same-sign lepton signature allows detection over a wider mass range (180 GeV up to 2 TeV) than does the missing- E_t signature. If squarks are heavier than gluinos, they will decay into gluinos, leading to the identical same-sign lepton signature (in this case, it may be difficult to tell the difference between direct gluino production and squark production followed by decay into gluinos). If the squark is lighter than the gluino, the same-sign lepton signature is expected to be much smaller.




FIG. 3-73. For events with two isolated same-sign leptons, the distribution of the invariant mass of the lower p_t lepton together with the two nearest jets chosen from the four jets with the highest p_t . Each lepton is required to have $p_t > 20$ GeV or alternatively one lepton must have $p_t > 40$ GeV and the other $p_t > 15$ GeV. The leptons must lie within $|\eta| < 2.5$. Events are required to have at least four jets with $p_t > 50$ GeV. The gluino mass was taken to be 300 GeV. The solid histogram is from gluino pair production and decay ($\tilde{g}\tilde{g} \rightarrow q\bar{q}\tilde{\chi}_1^+ q\bar{q}\tilde{\chi}_1^+ \rightarrow q\bar{q}\ell^+ \nu \tilde{\chi}_1^0 q\bar{q}\ell^+ \nu \tilde{\chi}_1^0$), whereas the dashed curve is due to $t\bar{t}$ production and decay ($t\bar{t} \rightarrow b\ell^+ \nu \bar{c}\ell^+ \nu$ jets, with $M_{\text{top}} = 150$ GeV). Unlike the signal, the dashed curve has not had any isolation cut applied. Such a cut would have eliminated it entirely.




FIG. 3-74. For events with two isolated same-sign leptons, the distribution of the invariant mass of the higher p_t lepton together with the two nearest jets chosen from the four jets with the highest p_t . Events are required to have one lepton with $p_t > 15$ GeV, one with $p_t > 65$ GeV, and at least four jets with $p_t > 50$ GeV. The leptons must lie within $|\eta| < 2.5$. The solid (dashed) curve is from $M = 300$ GeV ($M = 350$ GeV) gluino pair production and decay. The 300 GeV (350 GeV) gluino yields 6000 events (12,000 events) per year. (The cross section rises quickly with mass because of our cuts). The dashed curve has been divided by two for display purposes. This calculation was done using a parton-level Monte Carlo including resolution; a more realistic calculation will add tails to the peaks due to radiated jets that accidentally are near the lepton.

3.4.6. Direct production of charginos and neutralinos

In the previous sections, the production of charginos and neutralinos via the decay of heavy scalar quarks was described. In this section, we consider their direct pair production via Drell-Yan processes: $pp \rightarrow Z^*, \gamma^* \rightarrow \chi^\pm \chi^\mp$ and $pp \rightarrow W^* \rightarrow \chi^\pm \chi^0$. Here χ^\pm and χ^0 represent any of the several chargino and neutralino states present in the theory. In the minimal supersymmetric standard model there will be 4 χ^0 and 2 χ^\pm states. The χ 's undergo sequential decays via emission of W 's and Z until they reach the lightest state, which is assumed to be stable. This leads to sequences such as: $\chi_i^\pm \rightarrow W^\pm \chi_{i-1}^0$, $Z \chi_{i-1}^\pm$ and $\chi_i^0 \rightarrow W^\pm \chi_{i-1}^\mp$, $Z \chi_{i-1}^0$. The W 's and Z 's may be real or virtual, depending on the mass splittings.

While the rates for direct Drell-Yan production are much smaller than those for production via decay of scalar quarks, the events are cleaner if one searches for the signal in purely leptonic channels by requiring all of the vector bosons radiated during the decay to go into electrons or muons. In these channels, the hadronic multiplicity of the final state arises entirely from initial state radiation, and allows lepton isolation cuts to be rather efficient for the signal, while strongly suppressing backgrounds from semileptonic decays of heavy quarks. However, since t quarks will produce isolated dilepton pairs with high rate, we will concentrate on the case of three or more leptons. These multilepton final states arise from the sequential decays of the more massive χ 's. The details of the following analysis can be found in Ref. 63.

In Table 3-14 we give the values of the masses of the various states for some choices of the supersymmetric model parameters μ and M (we assume $\tan \beta = 2$). The parameters are chosen such that

the charginos and neutralinos are in a region beyond the reach of LEP II. This region is also complementary to the region where gluino and scalar-quark masses are relatively large and easily detectable with the SDC detector.

Table 3-14

Mass spectrum, in GeV, of the neutralino and chargino sector for various choices of the supersymmetric model parameters.

μ (GeV)	M (GeV)	χ_1^0	χ_2^0	χ_3^0	χ_4^0	χ_1^\pm	χ_2^\pm
-80	100	51	72	118	141	90	145
-80	200	72	104	108	225	94	224
-150	100	51	108	166	178	111	182
-150	200	96	144	172	225	155	226

To isolate the multileptonic signal we impose the following cuts: all of the leptons must be within 2.5 units of pseudorapidity, and the p_t of the two leading leptons should be larger than 20 GeV (electrons) or 15 GeV (muons). We selected two alternative cuts for the additional leptons: the p_t of the additional leptons should be larger than 10 GeV (cut *A*) or larger than the thresholds imposed on the leading leptons (cut *B*). For the isolation, we require no more than 2 GeV of energy within a cone of $\Delta R < 0.2$ around the lepton. To suppress the t -quark background we also require the total E_t of the event (after subtraction of the leptons) to be smaller than 30 GeV inside the region $|\eta| < 2.5$.

The resulting cross sections for the signal and the t -quark backgrounds are given in Table 3-15 for the cases of three and five leptons. The four-lepton signal is smaller because of suppressed couplings. Additional Standard Model backgrounds are negligible.

Table 3-15

Production rates (pb) for multileptonic final states from neutralinos and charginos. Cuts *A* and *B* are described in the text.

μ (GeV)	M (GeV)	3 ℓ (cut <i>A</i>)	3 ℓ (cut <i>B</i>)	5 ℓ (cut <i>A</i>)
-80	100	0.4	0.22	7.3×10^{-3}
-80	200	6.6×10^{-2}	3.4×10^{-2}	7×10^{-4}
-150	100	0.47	0.26	2×10^{-4}
-150	200	0.15	9.7×10^{-2}	2.1×10^{-3}
top bkgd	$M_{\text{top}} = 150$	0.8	0.26	$< 10^{-4}$
top bkgd	$M_{\text{top}} = 200$	4.8×10^{-2}	8.2×10^{-3}	$< 10^{-4}$

The Table shows that for the heavier t -quark mass, a large portion of the parameter space can be potentially covered in both the three and five lepton channels, while in the case of a lighter t quark, additional rejection power is required to see the signals.

3.4.7. Conclusions

We have shown that the SDC detector can readily isolate signals for a variety of supersymmetric particles such as gluinos, squarks, charginos, and neutralinos. In particular, there would be a significant signal for a 300 GeV gluino with only a fraction of the design luminosity via both the missing- E_t and the same-sign dilepton signatures. The detector-dependent background for the missing- E_t signature (mismeasured multijet events) can be reduced to a very small fraction of the signal. We have accounted for the major sources of mismeasurement, including a conservative estimate of nonGaussian resolution tails. The forward calorimeter plays an essential role in the elimination of missing- E_t backgrounds, and the required performance has been quantified. Furthermore, the SDC detector allows us to identify same-sign dilepton events and thereby provides a powerful tool for determining the gluino mass.

3.5. Heavy boson searches

New Z boson

Some models that enlarge the gauge group of the Standard Model predict the existence of new charged and neutral gauge bosons. The details of the masses and couplings are model dependent. We will concentrate on a new neutral gauge boson. Should such a boson be discovered, one would like to determine its mass, width, and couplings to quarks and leptons. Given a specific theoretical model all of these properties can be predicted. The new Z can be detected trivially via its decay to e^+e^- and $\mu^+\mu^-$. Observation of these modes can be used to determine the mass, width and the product of the production cross section (σ) and branching ratio (BR) to e^+e^- or $\mu^+\mu^-$. Most models respect lepton universality and hence the last two quantities are expected to be equal. The detection of the $\tau^+\tau^-$ mode was discussed in the EoI and LoI and will not be discussed further here.

By measuring the angular distribution of the leptons one can gain information on the helicity structure of the couplings of the new Z to quarks and leptons. The forward/backward asymmetry of the leptons is particularly useful. Events are selected for which the new Z is moving with pseudorapidity η . Since the large- x part of the quark distribution for a proton is larger than that for the antiquarks, for sufficiently large η the quark (anti-quark) that produced the Z is likely to have been moving in the same (opposite) direction as the Z itself. If the couplings of the quarks to the Z violate parity, the Z will then be produced with some preferred helicity. If the leptonic couplings also violate parity, then by determining the lepton sign one can determine an asymmetry

$$A = \frac{\int_0^{\pi/2} d\cos\theta \, dN/d\cos\theta - \int_{\pi/2}^{\pi} d\cos\theta \, dN/d\cos\theta}{\int_0^{\pi/2} d\cos\theta \, dN/d\cos\theta + \int_{\pi/2}^{\pi} d\cos\theta \, dN/d\cos\theta}$$

where θ is the angle between the direction of the positively charged lepton and the Z in the Z rest frame. It has recently been suggested [64] that measurements of the τ polarization via its decay to $\pi\nu$ also enable the helicity structure of the new Z couplings to be determined.

For definiteness, the couplings to quarks and leptons used in this section are those of a new Z defined in an E_6 model [65]. Such a model contains a mixing angle α that determines the properties of the new Z . For a fixed mass of the new Z , σ and BR are functions of α . If we assume that there are no decay channels involving new exotic particles then the measurement of $\sigma \cdot BR$ will severely constrain α . In the following we choose two particular models for a new Z . The first has $\cos\alpha = -0.6$, and is chosen to give a particularly large value for the asymmetry (-16%). The second is the Z_η model with $\tan\alpha = \sqrt{3/5}$, and is chosen because its width and its production cross section are particularly small. A new Z with the same couplings to quarks and leptons as the Standard Model Z has an appreciably larger production rate and width than a new Z from an E_6 model. The properties of a number of different models for a new Z are summarized in Table 3-16.

Table 3-16

A summary of the properties of a number of different models for a new Z boson (all in the context of a general E_6 model). A new Z with the Standard Model couplings to quarks and leptons is also included for reference. The widths are in GeV, and the production cross sections are in pb, including only events within $\pm\Gamma$ of the peak, and not including any efficiency or acceptance factors. The cross section includes the expected branching ratio to the e^+e^- final state, assuming that the new Z decays only to Standard Model particles.

Property	$\cos \alpha = -0.6$	Z_η	Z_ψ	Z_χ	SM Couplings
$\Gamma(M = 800 \text{ GeV})$	8.5	5.0	4.2	9.2	21.4
$\Gamma(M = 4000 \text{ GeV})$	42.3	25.2	21.0	46.2	106.9
$\sigma(M = 800 \text{ GeV})$	2.1	1.2	1.1	2.4	4.3
$\sigma(M = 4000 \text{ GeV})$	0.004	0.0032	0.0027	0.0051	0.010




FIG. 3-75. The cross section $d\sigma/dM$ for the production of an $\ell^+\ell^-$ pair for perfect resolution (dotted), the SDC resolution for electron pairs (solid), and the SDC resolution for muon pairs (dashed), as a function of the lepton pair invariant mass. The background is from the continuum production of lepton pairs (Drell-Yan). (a) The peak corresponds to a Z_η with a mass of 800 GeV. (b) The peak corresponds to a new Z with $\cos \alpha = -0.6$ and a mass of 800 GeV.




FIG. 3-76. As Fig. 3-75, except that the mass of the new Z is taken to be 4000 GeV.




FIG. 3-77. The cross section $d\sigma/d\cos\theta$ for the production of a lepton pair. The expected SSC result for e^+e^- (solid curve) and $\mu^+\mu^-$ (dashed curve) is shown. A perfect detector, which neglects acceptance and resolution smearing, is shown for reference (dotted curve). The reconstructed dilepton invariant mass is required to be between 700 and 900 GeV. The longitudinal momentum of the dilepton pair is required to be greater than 500 GeV. Events appear in the plot if the total charge of the lepton pair as determined by the detector is zero. (a) for a Z_η with mass of 800 GeV. (b) for a new Z with $\cos\alpha = -0.6$ and a mass of 800 GeV.

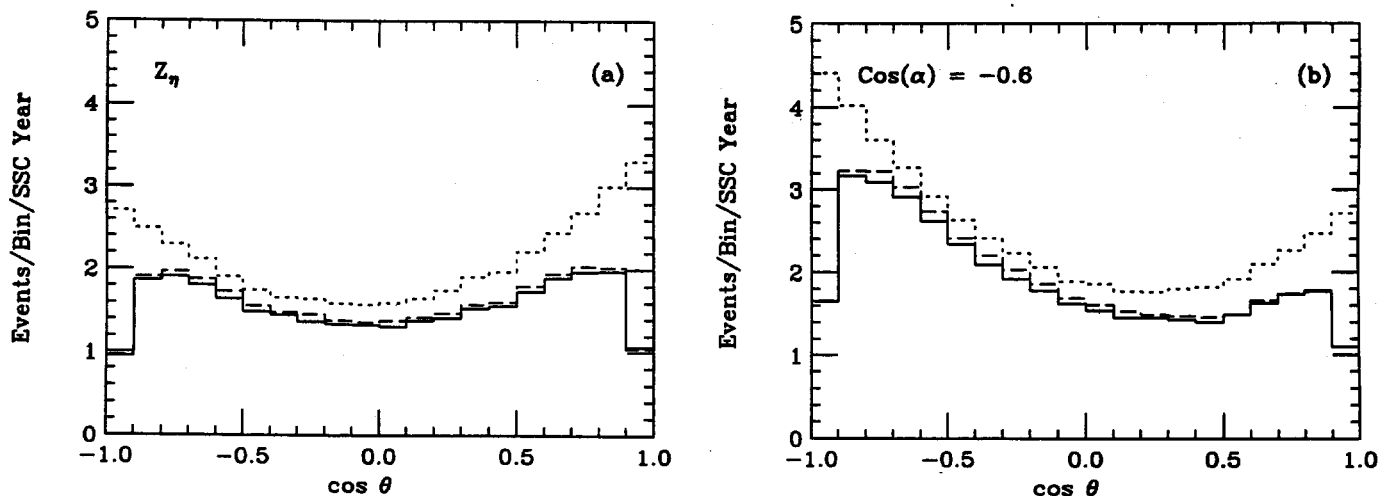


FIG. 3-78. As Fig. 3-77, except that a new Z with a mass of 4000 GeV has been assumed. The invariant mass of the dilepton pair was required to be between 3000 and 4000 GeV, and the longitudinal momentum was required to be greater than 1000 GeV.

In the e^+e^- and $\mu^+\mu^-$ channels there is no significant background except that arising from the Drell-Yan process. Figure 3-75 (a) shows the invariant mass spectra of dilepton pairs from the Z_η model as well as the background from continuum Drell-Yan processes. Figure 3-75 (b) shows the same distribution for the $\cos\alpha = -0.6$ model. The total number of events in the e^+e^- channel with invariant mass between 780 and 820 GeV is 7670 for the Z_η model, and 14570 for the $\cos\alpha = -0.6$ model. There is some uncertainty in the theoretical prediction of the production rates in a given model due to uncertainties in the quark distributions. Once other processes have been measured at the SSC these uncertainties should result in an error of less than 10% on the predicted rate. Hence a measurement of the event rate would be sufficient to provide some discrimination between different models. This method will not work if the new Z can decay into exotic channels that have not been identified. In this case, the branching ratio to $\ell^+\ell^-$ will be reduced from its predicted value. To avoid this ambiguity, we also need to measure the width and

asymmetry of the new Z .

In the e^+e^- channel, assuming the SDC baseline calorimeter performance defined in Table 3-1, the mass resolution of the detector is comparable to the natural width of the new Z . Hence both the mass and width can be determined in this channel. Fitting to the data in Fig. 3-75 results in values of $M = 799.95 \pm 0.1$ GeV and $\Gamma = 4.9 \pm 0.2$ GeV (the actual width is 5.0 GeV) for the Z_η case and $\Gamma = 9.1 \pm 0.2$ GeV (the actual width is 8.5 GeV) for the $\cos\alpha = -0.6$ case. We emphasize that these errors are statistical only. There is an additional error coming from the uncertainty in the energy scale of the electromagnetic calorimeter. By using the known mass of the Standard Model Z to provide an *in situ* calibration, as CDF was able to do using their data [66], this error could be as small as 0.1%. We have not attempted to determine the systematic error on the width determination. In the muon mode, one can measure the mass with a statistical error of 0.5 GeV.

Figure 3-76 shows the reconstructed invariant masses for a new Z of mass 4 TeV. There are a sufficient number of events to measure the mass. However, there are not enough events at design luminosity to use the measured value of $\sigma \cdot BR$ to constrain the theoretical models. Fitting to the data in Fig. 3-76 results in values of $M = 4000 \pm 5$ GeV for the e^+e^- final state. In the muon mode, one can measure the mass with a statistical error of 150 GeV.

The angular distribution of the negatively charged lepton in the rest frame of the pair, plotted in Fig. 3-77 and Fig. 3-78, clearly shows an asymmetry. The curves for the perfect detector neglect both acceptance effects and measurement errors. The acceptance effect is dominant, especially for the 800 GeV case (recall the Z was required to have a large longitudinal boost), and charge mis-measurement errors are negligible. The asymmetry is slightly more pronounced in the muon channel since the forward muon system enables a better measurement of the signs of forward going leptons than does the inner tracking system. The fall off near $\cos\theta = \pm 1$ is due to the acceptance of the SDC detector. Notice that the asymmetry is much larger for the $\cos\alpha = -0.6$ model, indicating that this measurement will be a powerful tool to discriminate between models.

3.6. Compositeness

If quarks are made of more fundamental objects with a binding scale of order Λ , one can expect a 4-fermion interaction between quarks for momentum transfers below Λ , of the form [67]

$$\mathcal{L} = -\frac{g^2}{2\Lambda^2}(\bar{u}_L\gamma^\mu u_L + \bar{d}_L\gamma^\mu d_L)(\bar{u}_L\gamma_\mu u_L + \bar{d}_L\gamma_\mu d_L)$$

Here g is the coupling strength of the interaction responsible for binding the quark constituents. Conventionally, one chooses $g^2/4\pi = 1$. The inclusive jet cross section will receive a contribution from the interference of this term with that arising from gluon exchange. This contribution does not fall as rapidly with the transverse energy as QCD does, and hence quark compositeness will manifest itself as a flattening of the inclusive jet cross section, $d\sigma/dE_t$, at large values of the jet E_t . Since the QCD calculation of the cross section has a 20–50% normalization uncertainty due to higher order corrections, structure function dependence and fragmentation effects, one typically fits to the shape of the QCD spectrum while letting the normalization float. As a result, energy-independent systematic errors (which typically change the normalization but only affect the shape of the distribution in second order) do not significantly compromise a compositeness measurement. An alternative technique, which is less sensitive to the jet energy scale, can also be used. This involves reconstructing the scattering angle of the jet in the two-jet center of mass system ($\cos\theta^*$), which is not directly affected by the jet energy scale. Here, we concentrate on the compositeness measurement derived from the jet E_t distributions since it places more stringent constraints on the detector performance.

In the SDC EoI, the maximum observable compositeness scale was defined to be that value of Λ for which at least 100 additional events are produced in a region where the QCD prediction is less than half as large as the compositeness prediction. This corresponds to a scale of roughly $\Lambda = 30$ TeV for one year

of data at the SSC design luminosity. A more general study of the energy dependent contributions to the systematic error on the jet energy scale was performed for the document submitted to the PAC in Summer 1990. The analysis presented at that time assumed that the non-linearity in the jet response was zero below 2 TeV, and increased to 2% at 5 TeV. Such a non-linearity corresponded to a systematic error which limited the bounds that could be set on the compositeness scale to a maximum of $\Lambda \approx 25$ TeV. It was shown that, for a non-compensating hadron calorimeter, the largest source of error arises from uncertainties in the fragmentation model which change the momentum spectrum of the hadrons in a jet. In the present study, we include a more complete model of the non-linear energy response of the calorimeter, in order to better quantify the impact of this factor on our ability to measure jet cross sections. In particular, we would like to assess whether the imperfect π/e response defined in Table 3-1 will limit the ability of the SDC detector to search for compositeness.

The single particle hadron response for the SDC baseline calorimeter design has been studied using CALOR89. The resolution and π/e response as a function of energy are those of Table 3-1. We note that preliminary test beam data indicates that somewhat better hadron calorimeter performance (both for resolution and for linearity) may be achieved in practice (further details appear in the Calorimeter section). Using the ISAJET Monte Carlo generator, along with a jet definition based on a cone of radius $R = 0.6$, the single particle non-linearity and resolution has been used to predict the induced jet non-linearity as a function of jet energy. There is a calibration scheme inherent in the definition of π/e response function described in Section 3.1.1. We assume that test beam data will be used to measure the π/e response below energies of several hundred GeV, and then extrapolated to higher energies. In addition, we have chosen to normalize the π/e response to be unity at an energy of 300 GeV. The resulting induced jet non-linearity is shown in Fig. 3-79. This figure displays the ratio of the jet E_t including the effects of single particle non-linearity and resolution to the jet E_t assuming perfect detector response (hence all other jet definition systematics are normalized out of this figure). Since π/e is less than one for energies below 300 GeV, the jet response is also less than one (it would eventually cross one and continue upwards for jet energies beyond 10 TeV, due to the finite energy at which the single particle π/e response is normalized to one).




FIG. 3-79. The non-linearity induced in the jet energy scale due to the non-linear single particle response of the SDC hadronic calorimeter.




FIG. 3-80. The fractional error on the jet energy scale induced by systematic errors in the knowledge of the single particle π/e response as a function of energy. The solid curve is for the worst case described in the text, namely perfect calibration of the response up to 100 GeV, and a 5% per TeV extrapolation error for higher energies. The dashed and dot-dashed curves correspond to an extrapolation error of 2% per TeV, and perfect calibration up to an energy of 500 (1000) GeV.

In principle, the only relevant feature of the non-linearity for a compositeness analysis is its uncertainty. If the jet energy scale were non-linear but perfectly known, it could be deconvoluted. We have used a naive model for the uncertainty—namely that we know the single particle π/e response perfectly up to some maximum energy, and then we assign an extrapolation error above that energy. Figure 3-80 indicates the fractional error in jet energy scale which is induced by several different models for the extrapolation error. The worst case (and the only case with observable consequences) is one in which we assume we know π/e up to an energy of 100 GeV, and then we make a 5% error per TeV in extrapolation (i.e., the single pion response is wrong by 5% at 1.1 TeV). Two other cases were considered where we calibrate to 500 (1000) GeV and then make an extrapolation error of 2% per TeV above that energy. In the worst case, an error of roughly 2% is induced in the jet energy scale at the highest attainable energies. This is almost precisely the non-linearity which was assumed in the SDC EoI, and leads to a systematic error limit on the search for compositeness of roughly $\Lambda = 25$ TeV. The other two curves represent more plausible models, and they induce a jet non-linearity which is below the 1% level at the highest energies encountered at the SSC. In conclusion, the performance of the SDC calorimeter, if properly calibrated, appears adequate to carry out compositeness searches at the SSC which will be limited only by statistics.

3.7. QCD tests

When the SSC turns on, we expect that our initial physics priorities will include measurements of Standard Model processes with large cross sections. These measurements will be used to study the detector and to verify the reliability of the theoretical predictions at these new high energies. This is particularly important for those processes that provide potential backgrounds to more exotic phenomena. In addition, one can imagine performing precise inclusive measurements, such as vector boson or heavy quark production rates, profiting from the looser trigger thresholds allowed during initial low luminosity running. Several unknowns involved in predicting rates for exotic processes and their backgrounds, such as structure function behavior at small- x or properties of the initial state radiation, might be constrained by studies performed during this early stage.

In this section, we provide cross sections for some bench-mark processes: single-jet and multi-jet inclusive rates, heavy quark cross sections and distributions, and single and multiple electroweak boson rates. These processes will not be used here to test the detector performance, but only to provide a reference for the expected rates. Hence, detection efficiencies will only be included via simple acceptance or momentum cuts.

3.7.1. Jet production

Jet production represents the largest source of high- Q^2 events in hadronic collisions and is therefore one of the most serious backgrounds for any search for new physics. Properly evaluating these backgrounds will rest on our ability to reliably calculate jet cross sections. QCD estimates of jet production rates are now available, including the full next-to-leading order corrections (at $O(\alpha_s^3)$), for single-jet inclusive distributions[68], and at the leading relevant order in α_s for production of n -jet events, with $n \leq 5$ [69].

In the case of single-jet inclusive processes, the $O(\alpha_s^3)$ calculations bring considerable benefits. For the first time, they predict the dependence of the cross section on the algorithm used to define the jet. Furthermore, they reduce the theoretical uncertainty arising from the choice of the factorization scale Q^2 used in evaluating the coupling constant and the structure functions. There are uncertainties remaining, due to the parametrization chosen for the parton distribution functions, and to the extrapolation of these functions into the small- x region ($x < 10^{-3}$). New data from HERA will help to reduce these uncertainties before SSC turn-on, but direct measurements of the jet cross sections at SSC energies will be the ultimate test that the theory is indeed under control.

The measurement of jets at small p_t and large rapidities, in addition to testing the performance of the forward calorimeters that will be used for the missing- E_t and jet tagging analyses, might also shed some light on the properties of gluons at very small- x [70]. Rates in these regions are so large that perhaps only by running at low luminosity will reasonable measurements be possible.

In Fig. 3-81 we plot the single-jet inclusive production cross sections in different regions of the detector. The rates correspond to the production of jets with p_t larger than the given threshold p_t^{\min} (in GeV). For the forward region, $3 < |\eta| < 5$, we also provide rates as a function of the jet energy. At nominal SSC luminosity, the rate for central jets with $p_t > 1$ TeV is roughly 1 Hz, while in the forward region we expect jets with $p_t > 100$ GeV and $E > 10$ TeV to be produced at a 10 Hz rate.

In Fig. 3-82, we plot the invariant mass distribution for QCD multi-jet events. The cases of 2, 3, 4, and 5 jets are considered. The distributions were calculated using the PAPAGENO generator [71] which uses an approximation to the QCD matrix elements introduced in Ref. 72. A comparison between several available approximations (see Ref. 73) and the exact results can be found in Ref. 69. The jets are separated by $\Delta R > 0.7$, they lie within $|\eta| < 2.5$ and have $p_t > 200$ GeV. Notice that at large invariant masses the curves for 2, 3, and 4 jets merge together, due to the large phase-space available.




FIG. 3-81. The inclusive single-jet cross section for jets with $p_t > p_t^{\min}$ in various regions of the SDC detector.




FIG. 3-82. The cross section for multi-jet production. The invariant mass of the jets is plotted.

3.7.2. Heavy quark production

b quarks

Among the various processes which contribute to jet final states, heavy quark production plays a fundamental role. Bottom quarks are produced in abundance and offer the potential for interesting studies. The total production cross section has been recently calculated to next-to-leading order (at $O(\alpha_s^3)$) in perturbative QCD [74], and has shown reasonable agreement with current hadron collider data [75]. However, large higher order corrections are known to be present at SSC energies [76], and hence a measurement of the total b cross section would provide a fundamental test of QCD in the small- x region.

At the present moment, predictions of the total b cross section at SSC energy vary over at least one order of magnitude, from less than 0.1 to more than 1 mb. The actual result depends strongly on the parametrization chosen for the input gluon distributions and on the choice of factorization scale. Table 3-17 summarizes the inclusive b quark production cross section, evaluated at $O(\alpha_s^3)$, for $p_t(b) > 10$ GeV (the $p_t(b)$ cut removes at least some of the ambiguity related to the small- x region of the structure functions). The results were obtained by using different sets of parton distributions [40], different values for the quark mass m_b , and different factorization scales f , defined by $\mu^2 = f^2(p_t^2 + m_b^2)$.

Having fixed the values of $m_b = 5$ GeV and $f = 1$, Fig. 3-83 indicates the inclusive production cross section for b quarks with transverse momenta larger than the threshold given by p_t^{\min} (in GeV). Higher order corrections due to the bremsstrahlung of more than one gluon from the final state b are not included

Table 3-17

The inclusive production cross section (μb) for a single b quarks with $p_t > 10$ GeV. The factor f is defined in the text.

m (GeV)	f	DFLM160	DFLM360	HMRSB	MTB
4.8	0.5	73	260	101	73
4.8	2	65	177	89	66
5.2	0.5	68	240	93	67
5.2	2	61	163	82	61

in the plot, but can be calculated within perturbative QCD[77]: their effect is to soften the spectrum of the b . As the figure shows, cross sections are very large. The rapidity distribution of b quarks above a given p_t (20, 100, 200 and 500 GeV) is shown in Fig. 3-84. Both the rapidity and p_t distributions will not differ appreciably for the \bar{b} .

For the purpose of b studies where both b and \bar{b} have to be detected, such as mixing and CP violation, Table 3-18 contains a compilation of the production cross sections for b -pairs contained within the tracking region ($|\eta_b|, |\eta_{\bar{b}}| < 2.5$), as a function of the p_t threshold for either the softest (p_t^{\min}) or the stiffest (p_t^{\max}) momentum. The results have been derived using the full next-to-leading order ($O(\alpha_s^3)$) calculation[78], with HMRSB structure functions, $m_b = 5$ GeV and $f = 1$. Notice that requiring both quarks to have momenta larger than a given threshold implies a loss in rate by a factor between 5 and 10.




FIG. 3-83. The inclusive production cross section for a single b quark with $p_t > p_t^{\min}$.




FIG. 3-84. The inclusive pseudo-rapidity distribution for b quarks.

Table 3-18

The production cross section (μb) for b pairs with $|\eta(b, \bar{b})| < 2.5$ as a function of the p_t requirements on the individual b 's.

p_t^{cut} (GeV)	$\sigma(p_t^{\max} > p_t^{\text{cut}})$	$\sigma(p_t^{\min} > p_t^{\text{cut}})$
10	64	12
20	15	1.2
50	1.5	0.14

When the p_t of the b becomes large, the contribution of processes where the b pair comes from the splitting of a hard final state gluon become more important. Such pairs will tend to be highly correlated in phase space. To show this effect, Fig. 3-85 displays the azimuthal correlation, namely the distribution of the variable $\delta\phi = \phi(b) - \phi(\bar{b})$, as a function of different p_t thresholds. To simplify this comparison, we have rescaled the curve corresponding to $p_t^{\max} > 50$ GeV by a factor of 50, and the one corresponding to $p_t^{\min} > 50$ GeV by a factor of 500. A plot of the ΔR distribution shows the same behavior, with the peak arising from collinear production becoming increasingly important at larger values of p_t . In Fig. 3-86, we show the transverse momentum distribution of the b pair, corresponding to the momentum of the jet recoiling against the quarks. We required both b 's to lie within $|\eta| < 2.5$ and to have $p_t > 50$ GeV. The peak corresponds to the gluon splitting contribution, while the shoulder down to $p_t = 0$ comes from the back-to-back production.




FIG. 3-85. The distribution of the $b\bar{b}$ azimuthal correlation. Both b 's lie within $|\eta| < 2.5$ and satisfy the p_t requirement listed on the figure. We have rescaled the curve corresponding to $p_t^{\max} > 50$ GeV by a factor of 50, and the one corresponding to $p_t^{\min} > 50$ GeV by a factor of 500.




FIG. 3-86. The p_t distribution of a $b\bar{b}$ pair, with p_t of the softest b larger than 50 GeV.

In addition to the processes discussed previously, another significant source of heavy quarks is the splitting of soft gluons emitted in the evolution of a high energy jet. The heavy quarks produced in this way will be much softer than those produced by standard QCD matrix elements, and hence will not affect the inclusive p_t distributions discussed above. However, they produce a significant contamination inside ordinary jets, contamination that might be important for some background studies. These effects can be calculated within perturbative QCD in the leading-log-approximation[79]. In Fig. 3-87, we plot the average c and b pair multiplicity inside a gluon jet of the given energy[80]. For jet energies above 1 TeV, 10% of the jets have at least one b pair in them and 20% have at least one c pair. For reference, the momentum distribution of b 's inside a 500 GeV gluon jet, calculated using HERWIG[81], is plotted in Fig. 3-88 in the form of a fragmentation function.




FIG. 3-87. The heavy quark pair multiplicity inside gluon jets, computed in perturbative QCD using the leading-log-approximation.




FIG. 3-88. The distribution for 500 GeV gluon fragmentation into b quarks.

t quark

Whether or not the t quark is discovered before SSC turn-on, it is expected that SSC will be the first real t -factory. The potential of the SDC detector to observe and study the properties of the t quark is discussed in detail in Section 3.3. Here we will only present, for reference, cross sections for the t as a function of its mass. The predictions for t production rates [74] are more reliable than those for the b and the c because of the larger values of x involved. However, some uncertainty is still present, related to the choice of structure functions and factorization scales. We explore this uncertainty by showing the variations encountered when these parameters are allowed to vary over reasonable ranges.

In Fig. 3-40, we have plotted the total inclusive $t\bar{t}$ cross section as a function of the t quark mass, which was varied over the 100 to 500 GeV range. If a fourth generation of quarks were to be produced, their production cross section would also be as shown in this figure. The band corresponds to the envelope of the smallest and largest rates obtained by varying the factorization scale between 0.5 and 2 times the quark mass, and by using the following sets of structure functions [40]: DFLM160, DFLM360, HMRSB, and MTB. In the range $100 < M_{\text{top}} < 200$ GeV we find that between 10^8 and 10^9 events are produced per year, at nominal luminosity and ignoring branching ratios.

The method for measuring the t -quark mass that was described in Section 3.3.2 uses t quarks produced with large p_t in order to reduce the combinatoric background to the mass measurement in the $\ell + \text{jets}$ final state. We therefore plot in Fig. 3-89 the inclusive production cross section for t quarks above a given p_t threshold, where we have used a reference value of $M_{\text{top}} = 150$ GeV. We obtain on the order of 10^5 events per SSC year above 1 TeV, which means that approximately one jet out of hundred above this energy would be a t jet.

A large fraction of $t\bar{t}$ events will be accompanied by additional jets from initial state radiation. The presence of these jets could influence the jet-tagging selection needed to extract the signal of W pair production from W fusion processes. For reference, Fig. 3-90 shows the rapidity distribution of the additional jet in $t\bar{t}$ events, where jets were required to have p_t larger than 20, 50, 100 and 200 GeV. The curves were normalized to the total t cross section. The fraction of events with an additional jet was 88% ($p_t > 20$ GeV), 44% ($p_t > 50$ GeV), 21% ($p_t > 100$ GeV) and 7% ($p_t > 200$ GeV). The results have been derived using the full next-to-leading order calculation ($\mathcal{O}(\alpha_s^3)$) for $t\bar{t} + \text{jet}$ production [78], and agree relatively well with the results of the HERWIG parton-shower Monte Carlo which has been used in Section 3.2.3 to study jet tagging efficiencies.




FIG. 3-89. The production cross section for t quarks with $p_t > p_t^{\min}$.




FIG. 3-90. The rapidity distribution for additional jets in $t\bar{t}$ production for various p_t thresholds. The curves are normalized to provide the fraction of events relative to the total $t\bar{t}$ cross section.

3.7.3. Prompt photon production

Production of hard photons, even though suppressed with respect to jet production, is very large. Complete next-to-leading order estimates of the inclusive cross sections are available [82], including the evaluation of the effects of quark-photon isolation requirements. These calculations were used to constrain the gluon distribution function by fitting lower energy prompt photon data [83]. It is possible that measurements of inclusive photon production at SSC will provide additional information and independent checks of the parton distribution functions.

In Table 3-19, we display the value of the differential distribution $(d\sigma/dp_t d\eta)_{\eta=0}$ as a function of p_t for both inclusive and isolated photons [84]. The isolation criterion requires the absence of hadronic energy of greater than 15% of the photon energy within a cone of $\Delta R < 0.1$ around the photon direction. The calculation used the ABFOW structure functions [83] and a factorization scale $Q = p_t/4$.

Table 3-19
The differential cross section, $d\sigma/dp_t d\eta$, at $\eta = 0$ for prompt photon production (pb/GeV).

p_t (GeV)	50	100	200	500	1000	2000
Inclusive	4.1×10^2	27	1.5	2.6×10^{-2}	9.6×10^{-4}	2.6×10^{-5}
Isolated	3.0×10^2	20	1.1	2.0×10^{-2}	7.8×10^{-4}	2.2×10^{-5}

It has frequently been pointed out that the rapidity distribution of photons can discriminate between different models of the gluon structure function at small- x . In Fig. 3-91, we plot the photon rapidity distribution [84] at different values of p_t and for two different parametrizations of the input gluon distribution at $Q = 2$ GeV, $xG(x) = (1-x)^2/\sqrt{x}$ (curve A) and $xG(x) = (1-x)^4$ (curve B).

Another interesting process is the QCD production of γ pairs. It was shown in the Section 3.2.1 that this process constitutes a significant background to the observation of direct Higgs production in the 80 to 120 GeV mass region. In addition to being a background for Higgs searches, a measurement of the signal itself could be interesting as a test of QCD. Therefore, Fig. 3-92 shows the invariant mass distribution for photon pairs, as a function of the acceptance region. We required $p_t > 40$ GeV for both photons. The tree




FIG. 3-91. The photon rapidity distribution for different values of the photon p_t . Curves A and B are discussed in the text.




FIG. 3-92. The invariant mass distribution for photon pairs for several different rapidity cuts.

level and gluon fusion processes are included, whereas we neglected the bremsstrahlung diagrams [85] which provide a contribution which can be suppressed by isolation cuts.

3.7.4. W and Z production

The perturbative QCD calculations of the production of W and Z vector bosons are probably the most precise calculations available for hadronic collisions. The $O(\alpha_s)$ contributions have been known for some time [86], and the full $O(\alpha_s^2)$ results have recently become available [87]. The residual uncertainty coming from even higher order corrections, estimated by varying the renormalization and factorization scale over the range 10 to 10^3 GeV, is reduced to 10–15%. The uncertainty coming from the choice of structure functions is significantly larger. This precludes the possibility of using Z production as a 10% luminosity monitor. Nevertheless, the precise measurement of inclusive W and Z production can be a very useful probe of parton distributions. In Table 3-20, we summarize the total cross sections for W and Z production, evaluated in Ref. 87, for different selections of structure functions [40]. The sets HMRSE+ and HMRSE- refer to parametrizations of the input gluon distributions $xG(x)$, where Q_0^2 behaves like \sqrt{x} and $1/\sqrt{x}$ respectively in the $x \rightarrow 0$ limit. A complete analysis of the theoretical uncertainties in the calculation of these cross sections can be found in Ref. 87.

Table 3-20
The production cross sections (nb) for W and Z bosons.

	HMRSE +	HMRSE	HMRSE -	HMRSB	MTE	MTB
$W^+ + W^-$	90.6	157	313	236	284	278
Z	29.3	49.4	93.1	73.6	87.9	86.0

This table indicates that, at the nominal SSC luminosity, one expects several leptonic decays of the Z and several dozen leptonic decays of the W every second.

The inclusive p_t distributions for massive vector bosons are also available at next-to-leading order accuracy [88]. We plot them in Fig. 3-93, with the Z curve rescaled by 0.1*. The HMRSB structure

* These curves and the numbers appearing in the following Table were kindly provided by M.H. Reno.




FIG. 3-93. The inclusive W and Z p_t distributions ($d\sigma/dp_t$, in pb/100 GeV). The Z curve has been rescaled by a factor of 0.1 for clarity.

Table 3-21

The inclusive W and Z p_t distributions ($d\sigma/dp_t$, in pb/GeV). The factor f is used to study the Q^2 -scale dependence, and is defined by $Q^2 = f(p_t^2 + m_{Z,W}^2)$.

p_t (GeV)	f	50	100	200	500	1000	2000	5000
W : DFLM160	1	1.1×10^3	2.0×10^2	18	0.36	1.2×10^{-2}	2.7×10^{-4}	5.9×10^{-7}
W : DFLM260	1	1.5×10^3	2.5×10^2	21	0.38	1.2×10^{-2}	2.6×10^{-4}	5.2×10^{-7}
W : DFLM360	1	1.9×10^3	3.0×10^2	23	0.38	1.1×10^{-2}	2.4×10^{-4}	4.7×10^{-7}
W : HMRSB	0.5	1.2×10^3	2.1×10^2	19	0.38	1.3×10^{-2}	3.2×10^{-4}	8.6×10^{-7}
W : HMRSB	1	1.2×10^3	2.1×10^2	20	0.39	1.4×10^{-2}	3.3×10^{-4}	8.5×10^{-7}
W : HMRSB	2	1.2×10^3	2.1×10^2	19	0.38	1.3×10^{-2}	3.2×10^{-4}	8.6×10^{-7}
Z : DFLM260	1	4.0×10^2	78	7.8	0.16	5.7×10^{-3}	1.4×10^{-4}	4.3×10^{-7}
Z : HMRSB	1	5.3×10^2	10	9.2	0.16	5.3×10^{-3}	1.1×10^{-4}	2.7×10^{-7}

functions were used, with $Q^2 = p_t^2 + m_{Z,W}^2$. The results obtained using different structure functions or factorization scales are shown in Table 3-21.

Finally, the possibility of using large- p_t W or Z production to probe c and b structure functions has recently been discussed in Ref. 89, to which we refer for details.

Associated production of W and multi-jet systems is one of the main backgrounds to several interesting processes, such as t quark production. Exact tree-level calculations exist[90] for processes with up to 4 jets in the final state, but absolute rates can vary by as much as a factor of 2-3 because of the choice of Q^2 scale or, once again, structure functions. It is reasonable to expect additional theoretical progress in the next few years, such as the calculation of higher order corrections, but measurements of jet distributions in W plus multi-jet events will remain fundamental in establishing the background levels for other important analyses. The rates will depend on the precise cuts imposed on the events; as an example[90], Table 3-22 summarizes the rates for W plus 2, 3, and 4 jets resulting from the following requirements: (i) $p_t^{j,\ell} > 50$ GeV, (ii) $|\eta^{j,\ell}| < 3$, (iii) $\Delta R_{jj,\ell} > 0.4$, (iv) missing- E_t larger than 50 GeV. The symbol ℓ represents the charged lepton from the W decay. The leptonic branching ratio for one family is included in the rates. It can be seen that more than half a million lepton + 4 jet events are expected within one year of running.

Table 3-22

The production cross sections (pb) for $(W \rightarrow e\nu) + \text{jets}$. The cuts applied are described in the text.

# of jets	2	3	4
$\sigma(\text{pb})$	177	107	46

3.7.5. Multiple boson production

The large phase space available at the SSC makes it possible to produce events with several massive gauge bosons. These processes can be used as tests of the Standard Model gauge couplings [91], and need to be measured to confirm the estimates of backgrounds for signals such as Higgs production or supersymmetry multi-lepton decays.

Complete calculations to order α_s are available for the production of boson pairs [92-95]. We provide in Table 3-23 the cross sections for pairs of W and Z bosons, obtained using different sets of structure functions.

Table 3-23

The production cross sections (pb) for boson pairs at order α_s . The row labeled by $\Delta(ZW)$ refers to the difference between the ZW^+ and ZW^- cross sections.

	DFLM160	DFLM260	DFLM360	HMRSB
ZZ	35.8	43.6	50.8	33.9
ZW^+	76.1	90.8	104.7	75.5
ZW^-	55.6	65.9	75.3	57.8
$\Delta(ZW)$	20.5	24.9	29.4	17.7
W^+W^-	277	335	388	278

The uncertainties due to changes in the Q^2 scale are particularly small for these processes, of the order of a few per cent. Unfortunately, one expects this stability to disappear after the gluon fusion ($O(\alpha_s^2)$) contributions are added [96]. In the case of Z pairs and W pairs, this precludes the possibility of constraining the structure functions. However, in the case of the ZW processes, the contribution of the gluon fusion processes can be cancelled by measuring the difference between ZW^+ and ZW^- . The theoretical estimate for this difference is therefore extremely stable, and provides an alternative powerful tool to probe QCD and the parton distributions. Including the leptonic branching ratios (1%) and a pessimistic overall 20% detection efficiency for the three leptons, one is left with a difference of the order of 500 events, corresponding to a statistical error of about 5%, which is significantly smaller than the differences between different structure functions.

Additional interesting tests of the Standard Model might come from the observation of multiple ($n > 2$) boson pairs. We summarize the relevant total rates, calculated in Ref. 97, in Table 3-24. These rates were obtained by setting the Higgs mass equal to 0. The presence of a Standard Model Higgs with a mass in the range 200 to 300 GeV will increase most of these rates by a factor of 2-4. If the Higgs were discovered in this mass range, the measurement of these multi-boson processes would therefore become an important test of the Standard Model couplings, even though the small cross sections might require either extended running or the exploitation of decay channels with jets or neutrinos to enhance the branching fractions.

Table 3-24
The production rates (pb) for multiple gauge bosons.

<i>WWW</i>	<i>WWZ</i>	<i>WZZ</i>	<i>ZZZ</i>	<i>WWγ</i>	<i>WZγ</i>	<i>ZZγ</i>	<i>WWWW</i>	<i>ZZZZ</i>
0.5	0.43	0.11	0.04	0.23	0.08	0.04	6×10^{-3}	2×10^{-4}

References:

1. F. Barreiro *et al.*, Nucl. Instr. and Meth. **A292**, 259 (1990).
2. F. Abe *et al.*, Phys. Rev. **D44**, 29 (1991).
3. F. Abe *et al.*, "Measurement of the Isolated Prompt Photon Cross Section in $p\bar{p}$ Collisions at $\sqrt{s} = 1.8$ TeV", submitted to Phys. Rev. Lett.
4. H. M. Georgi *et al.*, Phys. Rev. Lett. **40**, 692 (1978).
5. R. N. Cahn and S. Dawson, Phys. Lett. **B136**, 196 (1986); and S. Petcov and D. R. T. Jones, Phys. Lett. **B84**, 660 (1979).
6. F. Abe *et al.*, Phys. Rev. Lett. **68**, 448 (1992).
7. U. Amaldi *et al.*, Phys. Rev. **D36**, 1385, (1987); P. Langacker Phys. Rev. Lett. **63**, 1920 (1989); and G. Costa *et al.*, Nucl. Phys. **B297**, 244 (1988).
8. Z. Kunszt and W. J. Stirling, "The Standard Model Higgs at LHC: Branching Ratios and Cross Sections" *Large Hadron Collider Workshop*, CERN Report 90-10, Vol. II, p. 428.
9. J. Boucrot *et al.*, "Search for Neutral Higgs at LEP 200", *ECFA Workshop on LEP 200*, CERN Report 87-08, p. 312 (1987).
10. C. Barter *et al.*, "Detection of $H^0 \rightarrow \gamma\gamma$ at the SSC," High Energy Physics in the 1990's, Snowmass 1988, p. 98.
11. F. Abe *et al.*, "Measurement of the Isolated Prompt Photon Cross Section in $p\bar{p}$ Collisions at $\sqrt{s} = 1.8$ TeV", submitted to Phys. Rev. Lett.
12. M. Mangano, Solenoidal Detector Collaboration Note SDC-90-00103 (1990).
13. R. Raitio and W. W. Wada, Phys. Rev. **D19**, 941 (1979); J. N. Ng and P. Zakarauskas, Phys. Rev. **D29**, 876 (1984); and Z. Kunszt, Nucl. Phys. **B247**, 339 (1984).
14. W. Marciano and F. Paige, Preprint BNL-45805 (1991); J. F. Gunion, Preprint UCD-91-2 (1991); Z. Kunszt, Z. Trocsanyi, and W. J. Stirling, Preprint DTP 91/40 (1991); and A. Ballestrero and E. Maina, Preprint DFTT 29/91 (1991).
15. Z. Kunszt, Z. Trocsanyi, and W. J. Stirling, Preprint DTP 91/40 (1991); and A. Ballestrero and E. Maina, Preprint DFTT 29/91 (1991).
16. R. K. Ellis and Z. Kunszt, Nucl. Phys. **B303**, 653 (1988).
17. H-U. Bengtsson and T. Sjostrand, Comp. Phys. Comm. **46**, 43 (1987).
18. T. Han and S. Willenbrock, Preprint Fermilab-Pub-91/70-T.
19. J. F. Gunion, G. Kane, and J. Wudka, Nucl. Phys. **B299**, 231 (1988).
20. J. J. van der Bij and E. W. N. Glover, Phys. Lett. **B206**, 701 (1988).
21. E. W. N. Glover and J. J. van der Bij, Nucl. Phys. **B321**, 561 (1989).

22. U. Baur and E. W. N. Glover, Nucl. Phys. **B347**, 12 (1990); and Phys. Rev. **D44**, 99 (1991).
23. R. W. Brown and K. O. Mikaelian, Phys. Rev. **D19**, 922 (1979).
24. R. N. Cahn *et al.*, in *Proc. of the 1987 Workshop on Detectors and Experimental Areas at the SSC*, Ed. M. Gilchriese (World Scientific).
25. "SDC Letter of Intent," Solenoidal Detector Collaboration Note SDC-90-00151 (1990).
26. R. D. Kauffman and C. P. Yuan, Phys. Rev. **D42**, 956 (1990);
G. L. Kane and C. P. Yuan, Phys. Rev. **D40**, 2231 (1989); Phys. Rev. **D40**, 2244 (1989); and
J. F. Gunion *et al.*, Phys. Rev. **D40**, 2223 (1989).
27. A. B. Wickland, " $Z \rightarrow \text{hadrons}$ ", Solenoidal Detector Collaboration Note SDC-91-00081 (1991).
28. U. Baur and E. W. N. Glover, Phys. Lett. **B252**, 683 (1990); and
V. Barger *et al.*, Phys. Rev. **D44**, 1426 (1991); and Phys. Rev. **D44**, 2701 (1991).
29. M. Mangano, P. Nason, and G. Ridolfi, Preprint IFUP-TH-32-91.
30. J. F. Gunion and L. H. Orr, Preprint UCD-91-15; and
J. F. Gunion *et al.*, Preprint UCD-91-29.
31. J. Ohnemus and J. F. Owens, Phys. Rev. **D43**, 3626 (1991).
32. M. L. Mangano, SSC-SDE-32.
33. M. S. Chanowitz and M. K. Gaillard, Nucl. Phys. **B261**, 379 (1985).
34. M. S. Chanowitz and M. Golden, Phys. Rev. Lett. **61**, 1053 (1988); and Phys. Rev. Lett. **63**, 466 (1989).
35. M. S. Berger and M. S. Chanowitz, Phys. Lett. **B263**, 509 (1991);
D. A. Dicus *et al.*, Preprint UCD-91-10 and Phys. Lett. **B258**, 475 (1991); and
V. Barger *et al.*, Phys. Rev. **D42**, 3052 (1990).
36. The OPAL Collaboration, Phys. Lett. **B236**, 364 (1990);
the ALEPH Collaboration, Phys. Lett. **B236**, 511 (1990); and
the DELPHI Collaboration, Phys. Lett. **B242**, 536 (1990).
37. R. Ansari *et al.*, CERN-PPE-91-162;
F. Abe *et al.*, Phys. Rev. **D44**, 29 (1991).
38. F. Abe *et al.*, Phys. Rev. Lett. **68**, 448 (1992).
39. P. Langacker, Preprint UPR-0492T (1992).
40. M. Diemoz *et al.*, Z. Phys. **C39**, 21 (1988);
A. Harriman *et al.*, Phys. Rev. **D42**, 798 (1990); and
J. Morfin and W. K. Tung, Z. Phys. **C52**, 13 (1991).
41. A. Barbaro-Galtieri *et al.*, "Measurement of the Top Mass in the $e\mu$ Channel," Solenoidal Detector Collaboration Note SDC-90-00133 (1990).
42. D. Decamp *et al.*, Phys. Lett. **B244**, 551 (1990).
43. B. Hubbard, "Fragmentation Properties of Jets Produced in Proton-Antiproton Collisions at $\sqrt{s} = 1.8$ TeV", PhD Thesis, LBL-27687
44. A.B.Wicklund, Solenoidal Detector Collaboration Note SDC-91-00051 (1991).
45. For a review of Higgs bosons, see J.F. Gunion, H.E. Haber, G. Kane, S. Dawson, *The Higgs Hunters Guide*, Addison Wesley (1990).
46. R.M. Barnett, J.F. Gunion, H.E. Haber, I. Hinchliffe, B. Hubbard, and H.-J. Trost, Solenoidal Detector Collaboration Note SDC-90-00141 (1990); and new paper in progress.
47. B. Hubbard, " B Jet Tagging with the SDC Tracking System," SDC note SSC-SDC-31 and Solenoidal Detector Collaboration Notes SDC-92-00212 (1992).

48. An ℓ in either $\ell^+-\ell^-$ or ℓ - τ events can also originate from a leptonic τ decay. This contribution is small because of the combination of the τ branching ratio and the high- p_t cut on the lepton. An ISAJET calculation predicts a correction to the predicted $\ell^+-\ell^-$ rate from WW of approximately 10% for leptons with $p_t > 40$ GeV, and this correction can be determined with high precision. The τ 's from H^+ decays are not an important source of ℓ 's, because $B(H \rightarrow \tau\nu)$ is small for $\tan\beta$ values where $B(t \rightarrow H^+b)$ is large.
49. E. Witten, Nucl. Phys. **B188**, 513 (1981);
S. Dimopoulos and H. Georgi, Nucl. Phys. **B193**, 150 (1981); and
N. Sakai, Z. Phys. **C11**, 153 (1981).
50. For reviews and other references see J. F. Gunion, H. E. Haber, G. Kane, and S. Dawson, *The Higgs Hunter's Guide* (Addison-Wesley Publishing Company, Redwood City, CA, 1990), and H. E. Haber, in *Review of Particle Properties*, Phys. Rev. **D45**, Part 2, (June 1992).
51. K. Inoue, A. Kakuto, H. Komatsu, and S. Takeshita, Prog. Theor. Phys. **68**, 927 (1982) (E: **70**, 330, (1983)); and **71**, 413 (1984); and
R. Flores and M. Sher, Ann. Phys. (NY) **148**, 95 (1983).
52. J.F. Gunion and H.E. Haber, Nucl. Phys. **B272**, 1 (1986).
53. P. Fayet, Phys. Lett. **69B**, 489 (1977); and
G. Farrar and P. Fayet, Phys. Lett. **76B**, 575 (1978).
54. J. Ellis *et al.*, Nucl. Phys. **B238**, 453 (1984).
55. See, e.g., S. Dimopoulos *et al.*, Phys. Rev. **D41**, 2099 (1990).
56. Explicit forms for the chargino and neutralino mass matrices can be found in Appendix A of Ref.[52].
57. R.M. Barnett, J.F. Gunion, and H.E. Haber, Phys. Rev. Lett. **60**, 401 (1988); and Phys. Rev. **D37**, 1892 (1988).
58. T. Hessing and S. Behrends, "Inclusive Jet Spectrum E_t Corrections and Resolution Unsmearing", Internal CDF Note CDF-1132 (1990).
59. R.M. Barnett, E. Carlson and I. Hinchliffe, SDC note SSC-SDE-10 (September 1989); also LBL-27797.
60. A. Bartl *et al.*, Z. Phys. **C52**, 477 (1991);
H. Baer *et al.*, to appear in *Proc. 1990 DPF Summer Study on High Energy Physics: "Research Directions for the Decade,"* Snowmass, CO, FSU-HEP-901110; and
G.G. Ross, F. Zwirner, *et al.*, in *Proc. Large Hadron Collider Workshop*, Aachen, 4-9 October 1990, edited by G. Jarlskog and D. Rein (CERN, 1990) vol. II, p. 605.
61. R.D. Field, A. White, and R.M. Barnett, Solenoidal Detector Collaboration Note SDC-91-00043 (1991).
62. R.M. Barnett, J.F. Gunion, and H.E. Haber, in *Proc. of 1988 Summer Study on High Energy Physics in the 1990's*, Snowmass, Colorado, June 27-July 15, 1988, edited by S. Jensen (World Scientific, Singapore, 1989), p. 230;
R.M. Barnett, J.F. Gunion, and H.E. Haber, to appear in *Proc. 1990 DPF Summer Study on High Energy Physics: "Research Directions for the Decade,"* Snowmass, CO, LBL-29768; and
H. Baer, X. Tata, and J. Woodside, FSU-HEP-910731 (1991) and Phys. Rev. **D41**, 906 (1990).
63. R. Barbieri *et al.*, Nucl. Phys. **B367**, 28 (1991).
64. M. H. Austern, J. D. Anderson, and R. N. Cahn, Preprint LBL-31858 (1992).
65. V. Barger *et al.*, 1986 Snowmass, pg. 224.
66. F. Abe *et al.*, Phys. Rev. Lett. **63**, 720 (1989).
67. E. Eichten, K. Lane, and M. Peskin, Phys. Rev. Lett. **50**, 811 (1983).

68. R.K. Ellis and J. Sexton, Nucl. Phys. **B269**, 445 (1986);
S. Ellis, Z. Kunszt and D. Soper, Phys. Rev. Lett. **64**, 2121 (1990); and
F. Aversa *et al.*, Phys. Rev. Lett. **65**, 401 (1990).
69. F.A. Berends and H. Kuijf, Nucl. Phys. **B353**, 59 (1991); and
F.A. Berends, W.T. Giele, and H. Kuijf, Phys. Lett. **B232**, 266 (1989).
70. L.V. Gribov, E.M. Levin, and M.G. Ryskin, Phys. Reports **100**, 1 (1983).
71. I. Hinchliffe, Papageno, unpublished.
72. Z. Kunszt and W.J. Stirling, Phys. Rev. **D37**, 2439 (1988).
73. C.J. Maxwell, Phys. Lett. **B192**, 190 (1987); Nucl. Phys. **B316**, 321 (1989); and
M. Mangano and S. Parke, Phys. Rev. **D39**, 758 (1989).
74. P. Nason, S. Dawson, and R.K. Ellis, Nucl. Phys. **B303**, 607 (1988); Nucl. Phys. **B327**, 49 (1988);
G. Altarelli *et al.*, Nucl. Phys. **B308**, 724 (1988);
W. Beenakker *et al.*, Phys. Rev. **D40**, 54 (1989); and
W. Beenakker *et al.*, Nucl. Phys. **B351**, 507 (1991).
75. C. Albajar *et al.*, UA1 Collab., Phys. Lett. **B256**, 121 (1991); and Phys. Lett. **B262**, 497 (1991).
76. J.C. Collins and R.K. Ellis, Nucl. Phys. **B360**, 3 (1991);
S. Catani, M. Ciafaloni, and F. Hautmann, Nucl. Phys. **B366**, 135 (1991); and
E.M. Levin, M.G. Ryskin, and Yu.M. Shabelskii, Phys. Lett. **B260**, 429 (1991).
77. B. Mele and P. Nason, Nucl. Phys. **B357**, 409 (1991).
78. M. Mangano, P. Nason and G. Ridolfi, to appear in Nucl. Phys. B.
79. A.H. Mueller and P. Nason, Nucl. Phys. **B266**, 265 (1986).
80. M. Mangano and P. Nason, to appear.
81. G. Marchesini and B. Webber, Nucl. Phys. **B310**, 461 (1988).
82. P. Aurenche *et al.*, Nucl. Phys. **B297**, 661 (1988);
H. Baer, J. Ohnemus, and J.F. Owens, Phys. Rev. **D42**, 61 (1990);
P. Aurenche, R. Baier, and M. Fontannaz, Phys. Rev. **D42**, 1440 (1990); and
E.L. Berger and J. Qiu, Phys. Rev. **D44**, 2002 (1991).
83. P. Aurenche *et al.*, Phys. Rev. **D39**, 3275 (1989).
84. P. Aurenche *et al.*, in *Proceedings of the LHC Workshop*, Aachen 4-9 October 1990, G. Jarlskog and
D. Rein eds., CERN 90-10, Volume II, p.69.
85. P. Aurenche *et al.*, Z. Phys. **C29**, 459 (1985).
86. G. Altarelli *et al.*, Nucl. Phys. **B246**, 12 (1984).
87. R. Hamberg, T. Matsuura, and W.L. van Neerven, Nucl. Phys. **B359**, 343 (1991); and
W.L. van Neerven and E.B. Zijlstra, Univ. of Leiden preprint, INLO-PUB-14/91 (1991).
88. P.B. Arnold and M.H. Reno, Nucl. Phys. **B319**, 37 (1989); and
R. Gonsalves, J. Pawlowski and C-F. Wai, Phys. Rev. **D40**, 2245 (1989).
89. C.S. Kim, Nucl. Phys. **B353**, 87 (1991).
90. F.A. Berends, W.T. Giele, and H. Kuijf, Nucl. Phys. **B321**, 39 (1989);
F.A. Berends *et al.*, Phys. Lett. **B224**, 237 (1989); and
F.A. Berends *et al.*, Nucl. Phys. **B357**, 32 (1991).
91. U. Baur and D. Zeppenfeld, Nucl. Phys. **B308**, 127 (1988).
92. J. Smith, D. Thomas, and W.L. van Neerven, Z. Phys. **C44**, 267 (1989).
93. B. Mele, P. Nason, and G. Ridolfi, Nucl. Phys. **B357**, 409 (1991); and
J. Ohnemus and J.F. Owens, Phys. Rev. **D43**, 3626 (1991)

- 94. J. Ohnemus, *Phys. Rev.* **D44**, 3477 (1991); and
S. Frixione, P. Nason, and G. Ridolfi, Univ. of Genova Preprint, GEF-Th-2 (1992), UPRF-92-323.
- 95. J. Ohnemus, *Phys. Rev.* **D44**, 1403 (1991).
- 96. E.W.N. Glover and J.J. van der Bij, *Nucl. Phys.* **B321**, 561 (1989).
- 97. M. Golden and S.R. Sharpe, *Nucl. Phys.* **B261**, 217 (1985);
V. Barger and T. Han, *Phys. Lett.* **B212**, 117 (1988); and
V. Barger, T. Han, and H. Pi, *Phys. Rev.* **D41**, 824 (1990).

4. Central tracking system

4.1. Introduction

The physics requirements for the SDC tracking system have been analyzed and carefully quantified in Ref. 1 and are briefly discussed in the next section. Table 4-1 concisely summarizes these physics requirements. The tracking system discussed in this technical proposal has been designed to meet these requirements at a minimum cost, within the constraints of our assessment of the technological maturity and capabilities of the various options. The resulting SDC tracking system features a large instrumented volume with a radius of 1.7 m, half-length of 4.3 m, and filled with a 2 T solenoidal magnetic field. The large value of BL^2 as well as capabilities for track vertexing and pattern recognition are key features needed to meet the physics goals.

4.1.1. Summary of physics requirements for tracking

The central tracking system contributes to the detection of isolated and nonisolated muons and electrons, low-multiplicity decays of heavy gauge bosons, τ 's, J/ψ 's, etc., and (as a veto detector) photons. It allows identification of tracks from secondary vertices to help tag b -quark jets, and measures jet multiplicities. It gives a detailed picture of each event that may provide critical insight for the unexpected. It is called upon to provide momentum analyzed electron tracks for calibration of the electromagnetic calorimeter, and to provide a trigger sensitive to charged tracks above a well-defined p_t threshold.

Examples of potentially important physics processes which yield isolated high momentum tracks include rare reactions such as Higgs production and decay via the chain $H^0 \rightarrow Z^0 Z^0 \rightarrow \ell^+ \ell^- \ell'^+ \ell'^-$ or heavy gauge boson decay via $Z' \rightarrow \ell^+ \ell^-$. More complicated events include those involving t -quark production and decay, where the final state may include b -quark jets with detached vertices, isolated W -decay leptons, nonisolated b -decay leptons, and W^- and H^+ -decay jets.

The track detector needs to be sensitive to charged particles emitted at rapidities up to about 2.5 to provide adequate geometric acceptance for the processes of interest. High efficiency is required ($> 97\%$) for higher- p_t tracks that are well isolated from jets, as these are often decay products of heavy particles, some coming from rare processes or ones in which several partons are of interest. Examples of these include Higgs boson and top quark production. Isolated leptons are particularly important, so that the tracker must work reliably in conjunction with the calorimeters or muon detectors to identify and measure these. In the electron case this imposes a restriction on the amount of material introduced by the detectors, since bremsstrahlung by the electrons adversely affects the matching in both energy and trajectory with the calorimeters.

The momentum measurement precision impacts a broad range of physics analyses, including sharp reconstruction of the masses of intermediate states such as Z^0 bosons and sign-of-charge determination at the highest energies. To reconstruct the Z^0 mass within its natural width for a Z^0 produced at the typical p_t of 50 GeV/c requires the resolution (specified as fractional uncertainty at 1 TeV/c) to be better than 20%.

Besides the electron and muon, another important class of flavor-identified partons is the b quark. For effective b quark tagging the detector must resolve track impact-parameters near the beam axis at the level of tens of microns, and must have rather good efficiency for finding multiple tracks within an event even in quite crowded regions. Similar capabilities are needed for τ lepton identification.

Information from the tracking system is required for the trigger on electrons and muons, to permit the low (10 GeV/c) p_t thresholds necessary for efficient detection of isolated leptons. The design of the detectors and their readout electronics must provide for delivery of a Level 1 trigger signature for charged particles with p_t above a 10 GeV/c threshold that is precise within 10%, with good rejection of the copious lower momentum particles.

Table 4-1
Summary of requirements for tracking system.

-
-
- 1) Acceptance and p_t resolution:
 - (a) $|\eta|$ coverage at least out to $|\eta| = 2.5$ ($H^0 \rightarrow 4$ charged lepton geometrical efficiency $\gtrsim 60\%$ for $m_H \geq 200 \text{ GeV}/c^2$).
 - (b) Reconstructed (as opposed to parametric) vertex constrained momentum resolution for isolated charged tracks of $\sigma_{p_t}/p_t^2 < 20\% \text{ TeV}/c$ for $|\eta| \leq 1.8$, allowed to rise to $\sigma_{p_t}/p_t^2 \rightarrow 100\% \text{ TeV}/c$ as $|\eta| \rightarrow 2.5$.
 - 2) Reconstruction efficiency within this acceptance $\geq 97\%$ for isolated tracks having $p_t \geq 10 \text{ GeV}/c$, with ≤ 0.1 false tracks of $p_t \geq 10 \text{ GeV}/c$ per trigger. (Efficiency greater than 90% for detecting all four leptons from $H^0 \rightarrow 4$ charged leptons, exclusive of lepton identification— E/p and η, ϕ matching—and trigger cuts.) This requirement is specified for design luminosity, but with occupancy assumed twice that calculated by Monte Carlo for $pp \rightarrow X$.
 - 3) Reconstruction efficiency for same as in (2) $\geq 90\%$ at $10 \times$ design luminosity.
 - 4) b -hadron tagging efficiency for top studies with $125 \text{ GeV}/c^2 \leq M_{\text{top}} \leq 250 \text{ GeV}/c^2$:
 - (a) Reconstruction efficiency $\geq 80\%$ for tracks of $p_t > 5 \text{ GeV}/c$, with less than 10% fakes, within jets of p_t up to $100 \text{ GeV}/c$ for b tagging using leptons.
 - (b) b -hadron tagging efficiency $\geq 25\%$ with $\geq 90\%$ purity using detached vertices. Implies impact parameter resolution $\leq 20 \mu\text{m}$ for stiff tracks, $\leq 100 \mu\text{m}$ for $p_t = 1 \text{ GeV}/c$, and $\geq 85\%$ efficiency for finding tracks with $p_t > 1 \text{ GeV}/c$ within jets of p_t up to $100 \text{ GeV}/c$.
 - 5) Material $\lesssim 15\%X_0$, $\lesssim 7\%X_0$ inside 50 cm (average over $|\eta| \leq 2.5$). For efficiency of electron ID (E/p cut).
 - 6) Position resolution at the calorimeter shower maximum detector of $\leq 5 \text{ mm}$ in $r\phi$ (where bremsstrahlung smearing occurs) and $\leq 2.5 \text{ mm}$ in z .
 - 7) Alignment relative to the muon system of $100 \mu\text{m}$ in $r\phi$ and 2 mm in z .
 - 8) Jet charged multiplicity measurement within 15% for jets up to $p_t = 500 \text{ GeV}/c$ (to distinguish isolated W 's, and to study fragmentation for QCD studies and background modelling).
 - 9) Resolution for measurement of the z component of the vertex of 2 mm , to separate pileup interactions.
 - 10) First-level trigger with momentum resolution $\sigma_{p_t}/p_t^2 \sim 10 (\text{TeV}/c)^{-1}$ —implies a 10% error for a $10 \text{ GeV}/c$ lepton.
 - 11) First-level trigger efficiency $\geq 96\%$ per track, with ≤ 0.05 false triggers per calorimeter trigger ϕ bin per crossing, over the range $|\eta| \leq 2.5$.
 - 12) Second-level trigger with momentum resolution $\sigma_{p_t}/p_t^2 = 5 (\text{TeV}/c)^{-1}$. Gives a 20% error for a $40 \text{ GeV}/c$ lepton for triggering on $Z \rightarrow \ell^+\ell^-$, $W \rightarrow \ell\nu$.
 - 13) Discovery potential—hard to quantify. In general want maximum capabilities from detector. Based on history, highest priority (other than isolated lepton of Higgs case above) would be reconstruction and impact parameter measurement of leptons within jets up to the largest jet p_t possible (at least $\geq 500 \text{ GeV}/c$). Desired reconstruction efficiency $\geq 50\%$.
 - 14) Survivability at standard \mathcal{L} for ≥ 10 years.
 - 15) Allows a natural path for upgrading to a system with survivability of ≥ 10 years at $10 \times$ standard \mathcal{L} with emphasis (e.g. momentum resolution, pattern recognition, isolated track efficiency) to be decided based on what is learned during initial running.
-
-

4.1.2. Summary of baseline detector configuration

The baseline tracking system which matches our physics goals is shown in Fig. 4-1. It is composed of successive layers of detection elements. The innermost are made of silicon strip detectors inside a radius of 50 cm and half-length of 2.6 m. These are configured as concentric cylinders in the central region, and layers of disks perpendicular to the beam in the forward and backward regions. For larger radius, the measuring elements are cylindrical superlayers of straw drift tubes for rapidities $|\eta| < 1.8$, and gas microstrip arrays, organized as disks, covering more forward rapidities. The cylindrical arrays measure as the primary coordinates (i.e., the one measured with highest precision), the azimuthal coordinates at given radii while the disk arrays measure as the primary coordinates the azimuthal coordinates at fixed z values. The emphasis on the azimuthal coordinate comes from the fact that this allows the most direct measurement of the bending in the solenoidal field and allows a direct use of the measurements in a first-level trigger organized to provide a cut based on a rough p_t determination.




FIG. 4-1. Full SDC tracking system including inner silicon layers, outer straw drift-tube superlayers and forward gas microstrip detectors.

The coordinates used to determine the polar angle are measured using small-angle stereo (about half of the measurements) with a position resolution of about 1–1.5 mm per silicon or microstrip layer or straw superlayer. This allows the measurement of the polar angle with an accuracy of ~ 1 mrad and the z position with an accuracy of ~ 1 mm at the origin or at the calorimeter face.

The transition from barrel to disk geometry avoids the buildup of material which would occur if a barrel geometry were continued out to very large z coordinates. This transition is different for the silicon than for the outer tracking devices, reflecting a higher priority for limiting the material in the silicon which sits closest to the origin, and the different cost dependence on length for the two outer tracking devices.

Silicon tracking system (STS)

A quarter section of the silicon system is shown in Fig. 4-2, and the subassembly dimensions are given in Table 4-2. The full silicon area is about 17 m^2 and the number of electronics channels about 6×10^6 . An option for the detector would be to have the layers inside of 15 cm (first two barrel layers) made of pixel detectors. The mechanical design effort for pixels is focussed on ladder structures and cooling schemes which will allow such an option without necessitating mechanical system changes. The mechanical design allows for replacement of the innermost layers in the case of upgrades or following performance degradation after several megarads of radiation dosage.




FIG. 4-2. Silicon tracker design.

Table 4-2
Dimensions for silicon tracker.

Barrel	r	z	Silicon area	
(1)	9 cm	30 cm	6.78 m ² for Barrel	
(2)	12 cm	30 cm		
(3)	18 cm	30 cm		
(4)	21 cm	30 cm		
(5)	24 cm	30 cm		
(6)	27 cm	30 cm		
(7)	33 cm	30 cm		
(8)	36 cm	30 cm		
Disks	r_{in}	r_{out}	z	Silicon area
(1)	15 cm	39 cm	33 cm	10.16 m ² for Disks (both sides)
(2)	15 cm	39 cm	38 cm	
(3)	15 cm	39 cm	44 cm	
(4)	15 cm	39 cm	52 cm	
(5)	15 cm	39 cm	61 cm	
(6)	15 cm	39 cm	72 cm	
(7)	15 cm	39 cm	85 cm	
(8)	15 cm	39 cm	102 cm	
(9)	15 cm	39 cm	122 cm	
(10)	22.5 cm	46.5 cm	146 cm	
(11)	28.5 cm	46.5 cm	182 cm	
(12)	34.5 cm	46.5 cm	218 cm	
(13)	40.5 cm	46.5 cm	258 cm	
Total area = 16.94 m ²				

The strips on the silicon detectors are chosen to have 50 μm pitch. This allows all bonding and circuitry to be located on one end of a detector, minimizing material and cost. Individual detector modules are 12 cm long, the maximum length allowing sufficient signal-to-noise, and are made of two 6 cm wafers bonded together, to minimize fabrication costs. The primary coordinate readout is the ϕ cylindrical coordinate for both barrel and forward planar detectors. A second set of strips at a small stereo angle on the backside of the detectors provides a coarser measurement of the second coordinate which is the z coordinate in the

barrel section and the r coordinate in the forward detectors. This allows a maximum contribution of the measurements on both sides of a given detector to the momentum determination. A choice of resolution of about 1.5 mm (stereo angle of 10 mr) for the coarsely measured coordinate is sufficient to avoid broadening the multiparticle mass resolution due to the angle determination in this coordinate. The information from an adjacent pair of detectors effectively provides a space point and track tangent for the pattern recognition. Figure 4-3 shows the arrangement of adjacent strip layers for the barrel.




FIG. 4-3. Orientation of double-sided strips on radially adjacent detector pairs.

The parameters of the system are determined by several general considerations. The pattern recognition role dictates the minimum number of detector layers. This issue is discussed in section 4.2; simulation studies indicate that seven layers averaged over η is a good choice. The very low occupancy of the silicon detectors allows excellent pattern recognition at luminosities much beyond the design luminosity, providing significant high luminosity capability for the SDC detector. The vertexing requirement dictates a ratio between outer and inner radii of about four in order for the vertexing accuracy to be comparable to the individual detector measuring accuracy. To allow extrapolation to the outer tracker with an accuracy ($\sigma = 2$ mm) comparable to the size of an outer tracker measuring element requires about 25 cm of radial extent for the silicon tracker. Our baseline design has an average of 27 cm of radial extent in both barrel and disk regions. In the very forward direction, the silicon is a major element determining the momentum measuring accuracy and thus should have the largest lever arm affordable. The rapidity coverage of the silicon is matched to that of the muon system.

Barrel outer tracking detector (OTD)

The barrel outer tracking system, shown in Fig. 4-4, is composed of five cylindrically concentric superlayers of straw drift tubes. Each superlayer is made of 6 or 8 layers of straw tubes (trigger superlayers have 8 layers, the others have 6). The straws are all 4 mm in diameter, which allows individual element occupancies to be sufficiently small. The straw tubes are arranged in modules of approximately trapezoidal cross section, each containing about 200 straws. Each straw module is essentially an independent tracking chamber with its own gas and power connections and its own electronics. The modules are described in more detail in section 4.4.3. The superlayers have straws running either parallel to the beam direction (axial superlayers) or at a small angle to this direction (stereo superlayers) in order to measure the coordinate (z) along the wire. The superlayers provide local track segments, characterized by an azimuthal angle and slope. The azimuthal offset in the stereo superlayers provides the z measurement.

All superlayers are divided in half at $z \approx 0$, with electronics located near the endcap calorimeters. Three axial superlayers are needed for a high- p_t track segment trigger, described in more detail in section 4.2.3, which uses two out of three segments to allow sufficient fake rejection and efficiency at high luminosity. The minimum number of stereo superlayers is two, one with wires running at about $+3^\circ$ to the beam direction, the other at -3° . One of them is placed between the outer two axial layers, and the other is




FIG. 4-4. Barrel outer tracking detector.

placed at smaller radius to improve the linking to the silicon inner tracker, which also provides polar angle information. The inner-most axial superlayer is placed at a radius close to half the maximum radius, where the occupancy is still small enough to allow good segment finding efficiency. Momentum measurement is accomplished by using both the inner and outer tracking systems in an integrated manner. The outer tracker improves the momentum resolution by a factor of ten over the silicon system alone, and improves the polar angle resolution by a factor of five. Table 4-3 summarizes the component information for the outer tracking system. There are a total of 1.37×10^5 straws.

Table 4-3
Barrel outer tracking system configuration.

Superlayer	Mean radius (m)	Channel count	Modules	Layers/ superlayer	z_{max} (m)	Stereo angle (°)
1	0.816	19,504	92	8 (trigger)	2.410	0
2	1.103	19,716	124	6	3.281	+3
3	1.351	32,224	152	8 (trigger)	4.033	0
4	1.488	26,712	168	6	4.033	-3
5	1.631	39,008	184	8 (trigger)	4.033	0

Intermediate angle tracking detector (ITD)

The ITD consists of tracking detectors positioned symmetrically on both sides of the interaction point between the inner silicon tracker and the endcap calorimeters. The maximum η coverage of the ITD is driven by the requirement to provide tracking information up to the maximum useful η coverage of the electromagnetic calorimeter, namely 2.8 units. This is somewhat larger than the η coverage of the silicon. The exact placement of the boundary between the barrel and intermediate regions depends on considerations of continuity of coverage, cost and the engineering of supports and access. Between the barrel and intermediate angle trackers, space must be available for support structures, terminations of the active regions of both trackers and passage of utilities and signals. We have chosen the intermediate angle region to start at an $|\eta|$ of 1.8. The detector (one side only) is shown in Fig. 4-5.




FIG. 4-5. Intermediate tracking detector (one side only).

The ITD consists, on each end, of three superlayers of microstrip tiles arranged projectively. This projective scheme provides direct measurement of ϕ vs. z in η bands which allows direct use of the measured information in the trigger. It provides good resolution and very low occupancy. The use of a three superlayer trigger reduces the combinatorial background to an acceptable level while providing sufficient resolution.

The minimum requirement from each superlayer is a space point on the track. This can be satisfied if each superlayer contains at least one layer of radial anodes (ϕ measurement) and one layer of small angle stereo anodes. Since one of the prime requirements of the ITD is to provide a Level 1 trigger on high- p_t tracks, it must be highly efficient at detecting tracks, at least in the radial anode layers used in the trigger. For this reason we choose to use an OR of two radial anode layers in each superlayer. We have also chosen two small angle stereo layers, of opposite inclination, to eliminate the confusion, ghosting, that occurs when only one small angle stereo layer is used. The four layers will be as closely spaced in z as possible, each superlayer occupying no more than 10 cm.

In order to measure the transverse momentum of a track, a disk-based ITD must measure both η and ϕ . Studies [2] have shown that for the first-level trigger, an acceptable p_t threshold turn-on can only be

List of supporting documents for baseline design

This proposal for the SDC baseline tracking system is based on several years of R&D and several very detailed conceptual design studies which provide supporting documentation. These supporting documents are: Requirements for the SDC Tracking System [1], Silicon Tracking Conceptual Design Report [3], SDC Modular Straw Outer Tracking System Conceptual Design Report [4], and Intermediate Angle Track Detector Conceptual Design Report [5]. In addition, SDC notes which are referenced in the appropriate sections address key issues for tracking, such as alignment and triggering.

4.1.3. Fiber tracker option

The baseline central tracking system has been chosen based on the present status of the various R&D efforts for tracking within the SDC. In the case of the outer tracker, very significant progress has been made in the two major options: tracking based on straw tubes and tracking based on scintillating fibers [6]. The former is less of an extrapolation from existing devices and has, therefore, been selected as the baseline choice. It is, therefore, the basis on which we have evaluated performance and have elaborated a reasonably complete conceptual design. It is the choice of OTD used in all discussions in sections 4.2 through 4.7.

An outer tracker based on scintillating fibers would also provide a powerful device and have some advantages, particularly in the case of luminosities significantly beyond the SSC design value. This technology is maturing rapidly and the SDC expects to be able to make a final choice before the end of the summer 1992. Key issues to be evaluated for both tracker options are system performance and, in the case of fibers, the production quality photon sensor. To allow evaluation of the status and potential of the fiber option, we have included a rather detailed discussion of a fiber OTD in section 4.8.

4.2. Rationale for design

The tracking system must meet stringent goals for vertexing, pattern recognition, momentum resolution and triggering. In the discussion in this report, we will assume about $12\ \mu\text{m}$ measuring accuracy per double-sided silicon detector ($17\ \mu\text{m}$ per side) and about $85\ \mu\text{m}$ for each outer multiple-element superlayer of straws. These errors include contributions from local alignment errors. With these numbers we can compare various alternatives, assuming perfect global alignment. For example, taking as a goal a momentum error of 20% at 1 TeV/c for the nonvertex constrained momentum measurement, we find that a system with 11 equally spaced silicon layers or 11 outer detector superlayers in our 2 T magnetic field would require an active tracking length of 80 cm and 2.2 m, respectively. The mixed system chosen by the SDC requires about 140 cm and therefore represents a compromise between limiting the radial extent and limiting the amount of silicon. Including the distance to the first measuring layer and leaving room for global alignment contributions to resolution, we arrive at the SDC tracking radius of 1.7 m. This allows most of the straw tube layers to be located at a radius $> 1\ \text{m}$ where the occupancy is fairly small. Indeed, even at several times the design luminosity the straw system provides reliable track segments for tracking. All track parameter measurements depend on the full tracking system to achieve the ultimate resolution desired.

With the SDC design, the silicon array provides a device with powerful pattern recognition capabilities and vertex tagging up to very high luminosity and over the full rapidity range. This results from the high precision of the device, its good two-track resolving power, its low occupancy (about 10^{-3} at the smallest radius at design luminosity) and enough measurements to significantly overconstrain the parameters of track candidates. At large rapidities the gas microstrips provide a low occupancy system with a geometry very well matched to that of the silicon disk arrays.

Similarly, the organization of the straw detector measurements into superlayers allows a substantial reduction of measurement combinations that must be investigated in the track finding. A superlayer track segment gives direction within about 0.01 radians, as well as position. This information can be used to eliminate many fragments of loopers and to sort segments from potential tracks into bins related to curvature as well as ϕ for the formation of tracks. The curvature information is independent of dip angle

even in the stereo layers, greatly facilitating the incorporation of these measurements during the pattern recognition.

Responsibility for providing a Level 1 trigger lies with the OTD and ITD and the geometry of these devices has been chosen with this trigger in mind.

4.2.1. Matching of detector to physics requirements

In this section we evaluate the proposed design in terms of the physics requirements. We consider the nominal parametric precision and some results of detailed simulation of several key physics processes. We address track finding efficiency and resolution, vertexing capability, matching to the outer detectors and particle identification, and triggering.

Parametric performance

The ability of the tracking system to measure a charged particle trajectory may be characterized by the (approximately Gaussian) errors on the five parameters that describe the helical trajectory. These parameters are taken to be the azimuthal angle ϕ_0 , cotangent of the polar angle of the track at its point of origin $\cot(\theta)$ (also written as $\tan(\lambda)$ where λ is called the dip angle) the transverse momentum p_t , the impact parameter relative to the beam line b , and the axial position z_0 at the point of the helix's closest approach to the beam axis.

We determine the expected parameter errors by constructing the error matrix for a set of measurements along a trajectory of fixed pseudo-rapidity through the model detector described by the geometrical data from Tables 4-2, 3 and 4. Measurement errors are $17 \mu\text{m}$ for each single-sided silicon measurement, $85 \mu\text{m}$ for each barrel straw superlayer, and $100 \mu\text{m}$ for each ITD tile. These values include the estimated contributions from local alignment errors (and individual layer inefficiencies in the case of drift tubes). The p_t error has been evaluated both with and without a pseudo measurement at the beam axis (beam constrained, BC) with $20 \mu\text{m}$ precision.

Table 4-5
Parametric resolutions for the barrel tracker at $\eta = 0$.

Track parameter	σ
$\Delta p_t/p_t$ at 1 TeV/c (BC)	0.15
$\Delta p_t/p_t$ at 1 TeV/c (No BC)	0.17
$\Delta \phi_0$ (mrad)	0.066
$\Delta \cot \theta$	0.0013
Δb (μm)	13
Δz_0 (mm)	0.77

For tracks of large enough momenta that multiple scattering and energy loss are negligible, the resolutions on track parameters at $\eta = 0$ are listed in Table 4-5. The parameter errors are plotted as a function of $|\eta|$ and p_t in Fig. 4-6. The effect of multiple scattering is included in the calculations made to produce these plots.

Because of the substantial difference in resolutions between the silicon and wire detectors, the transverse momentum determination is dominated by the position and direction of the silicon portion of the track combined with the wire measurement near the coil. The dip angle determination receives significant contributions from stereo measurements in both detectors.

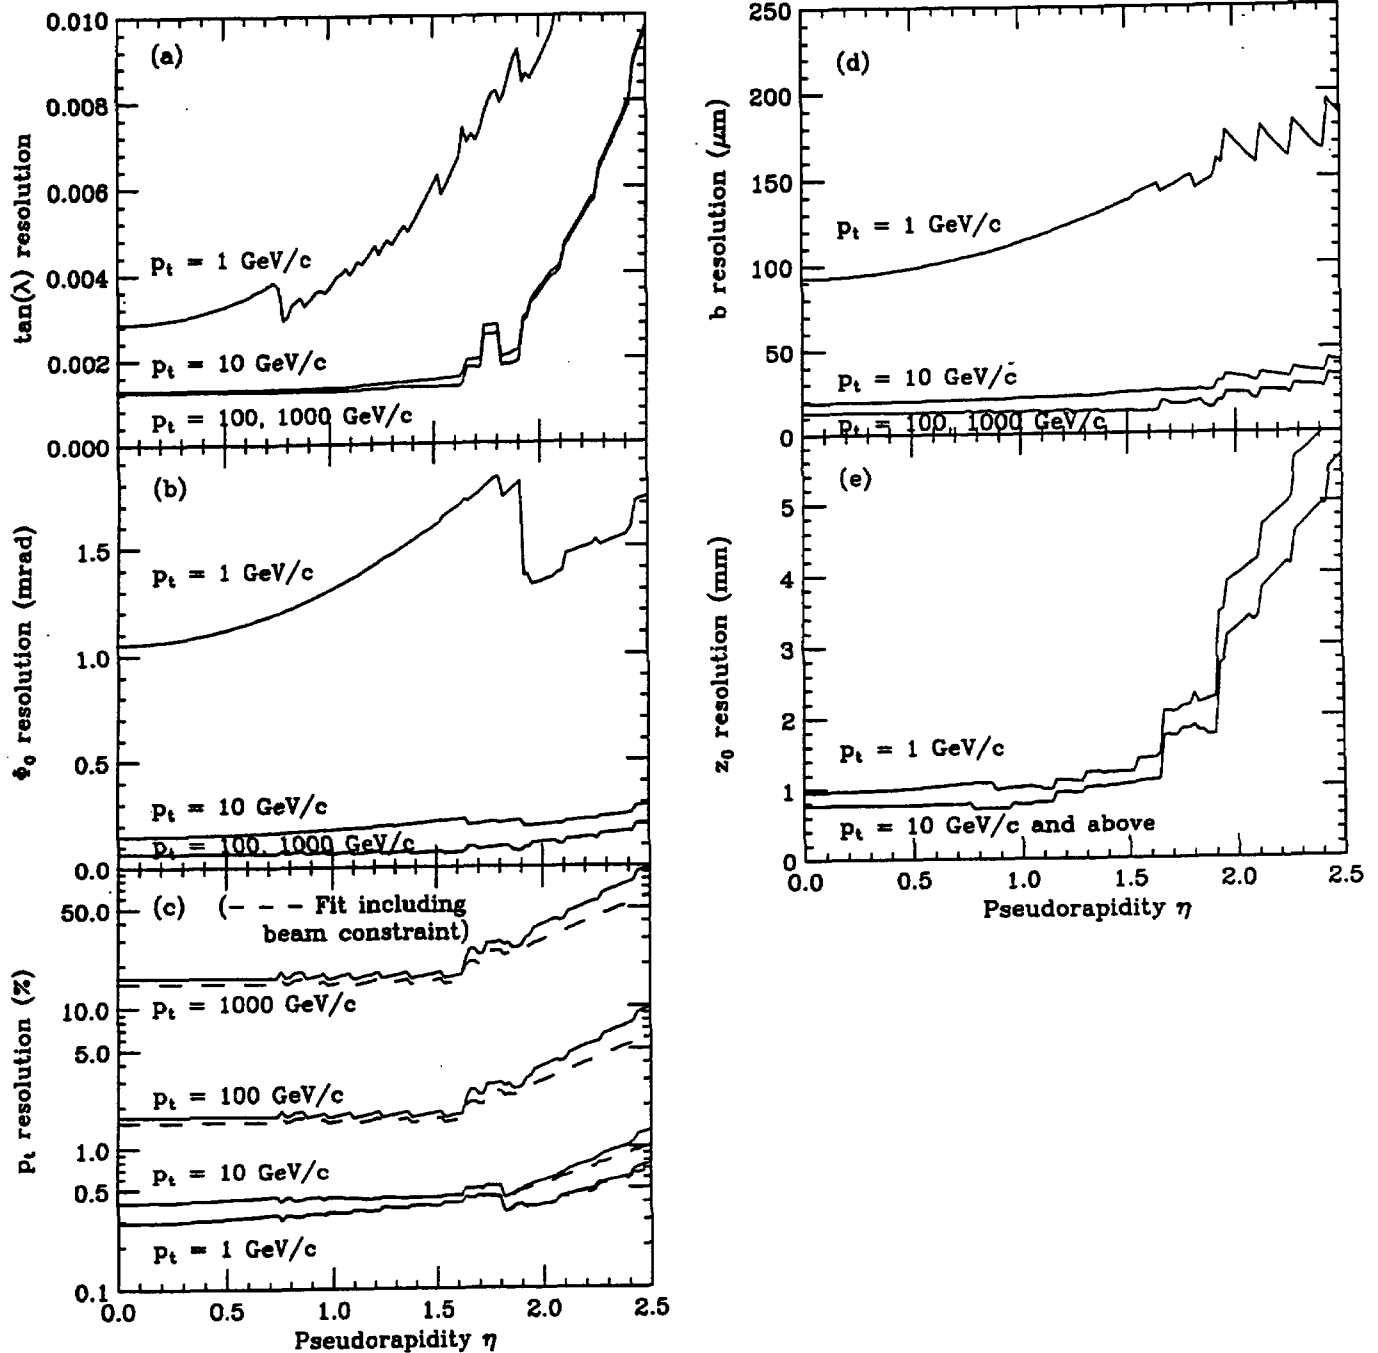


FIG. 4-6. Track parameter errors vs. rapidity for several values of p_t . Parameters are: a) tan of dip angle of the track at its point of origin, b) azimuthal angle ϕ_0 , c) transverse momentum p_t , d) impact parameter b , and e) axial position z_0 .

As a stand-alone system, the STS should provide an impact parameter measurement with an error small enough for B and τ tagging, using reasonably high- p_t tracks. This capability will allow good vertexing even in complex events, where the outer tracking efficiency could be degraded. Assuming n equally spaced double-sided silicon layers with fixed measuring accuracy σ_m , the error on the impact parameter b , σ_b , is just a function of the ratio of the outer to the inner radii of the silicon detector for high- p_t tracks, where multiple scattering can be neglected. That is, $\sigma_b = \sigma_m f(R_{Si}^{Out}/R_{Si}^{In})$. Taking for definiteness, $n = 8$, Fig. 4-7

shows σ_b vs. the ratio of radii for $\sigma_m = 12 \mu\text{m}$. To get down to $\sigma_b \lesssim 25 \mu\text{m}$ requires a ratio of radii $\gtrsim 4$. To minimize the effects of multiple scattering both $R_{\text{Si}}^{\text{In}}$ and the material associated with the innermost layers must be minimized. A reasonable choice, given the severe radiation environment close to the beam, is $R_{\text{Si}}^{\text{In}} = 9 \text{ cm}$ and $R_{\text{Si}}^{\text{Out}} = 36 \text{ cm}$. These dimensions provide the desired 27 cm radial coverage discussed earlier.




FIG. 4-7. Impact parameter error vs. detector geometry for the STS alone. Calculation is for 8 layers, as in central rapidity region.

For stiff tracks with full tracking information the impact parameter error will be improved by a factor of two compared to using only the silicon, because of the more precise curvature determination. The multiple scattering dominates the error for $p_t < 10 \text{ GeV}/c$. In this regime the outer tracker does not improve σ_b very significantly.

The parametric errors given here demonstrate the ability of the proposed tracking system to meet or exceed the SDC requirements, assuming that the design can be realized with performance close to nominal.

Simulation studies

The SDC tracking simulation allows for a flexible description of the detector in the GEANT framework. The detectors and inert material are described by a hierarchical volume structure which includes the inactive regions at the edges of detector elements. The GEANT routines transport particles through the apparatus and generate decays and interactions with material, including conversions, bremsstrahlung, and delta rays. Using this program and track reconstruction algorithms, we can evaluate the physics performance of the detector both at design luminosity and at higher luminosities. Some performance examples for key physics processes are discussed below.

Performance in Higgs events for various luminosities

The decays of neutral Higgs bosons, (mass = 300 GeV) $H^0 \rightarrow Z^0 Z^0 \rightarrow e^+ e^- \mu^+ \mu^-$ and $\mu^+ \mu^- \mu^+ \mu^-$, have been studied in some detail. The ISAJET generator is used to produce the Higgs event, and PYTHIA minimum bias events are added as background. The design luminosity is simulated by a Poisson-distributed mean of 1.6 background events for each bunch crossing, with a total of 7 crossings (-4 to $+2$ from the crossing with the Higgs) used to account for loopers and allow for drift times in the straw system. Only events where all four final state leptons had $|\eta| \leq 2.5$ are kept for further analysis. Figure 4-8(a)-(c) shows a view of a typical Higgs event (in this case decaying to four muons) as seen by the barrel detectors. The silicon information displayed shows the segments made by pairing neighboring layers. Raw hits in the

silicon are shown in b). The silicon resolution is much smaller than the size of a hit in the figure. The reconstruction of segments in the straw layers is based on the drift-time pattern for the full superlayer. Examples of the segment reconstruction in a region with many hits is shown in c).

Figure 4-9 shows the efficiency for track reconstruction vs. $|\eta|$ for all tracks with transverse momenta above 1 GeV/c in 200 Higgs events at design luminosity. For tracks with $p_t > 10$ GeV/c, the efficiency (averaged over $|\eta| \leq 2.5$) is 99%. The number of fake tracks (tracks which have more than 2 hits from another Monte Carlo track) with $p_t > 5$ GeV/c is found to be 0.03 ± 0.01 per event.

For this simulation, the lepton identifications are just based on using the Monte Carlo information in place of the calorimeter or muon system. Leptons are required to have $p_t > 10$ GeV/c and electrons to have $P_{\text{gen}}/P_{\text{fit}}$ within the range 0.7 to 1.4 to simulate the calorimeter energy and track momentum matching cut. In the case of electrons, the tracking momenta for these cuts are based on refitted values using only the silicon information in order to minimize the effects of bremsstrahlung. This results in a 6% higher efficiency for this cut, per electron, than using the full system. Note that all leptons passing the above cuts are used in the analysis; therefore such leptons from b decays and conversions are considered.

To pair the leptons in order to make Z^0 candidates, the following algorithm is used. All pairs of same type, opposite sign leptons are made and the invariant mass calculated. The invariant mass of the di-electrons is formed using the direction information from tracking but smearing the generated energies by the calorimeter resolution, since this typically provides the most accurate value of the electron energies. The di-muon invariant mass is calculated using the reconstructed momentum from the central tracking system only. The list of Z^0 candidates is made by ordering the pairs by their closeness to the Z^0 mass and then accepting pairs so that each lepton is used only once. The invariant masses of the electron pairs are shown in Fig. 4-10(a), and for the muon pairs in Fig. 4-10b. A cut requiring the invariant mass to be within 10 GeV of the Z^0 mass is then used to further restrict the list before the Higgs is reconstructed.

The four-lepton invariant mass distribution is plotted in Fig. 4-10(c) for events with two Z^0 candidates. Figures 4-10(a)–(c) are for the design luminosity of $10^{33} \text{ cm}^{-2}\text{s}^{-1}$ for which the efficiency for reconstructing the Higgs is 84%. This includes a small loss for events where the leptons do not all pass the 10 GeV/c p_t cut. Table 4-6 summarizes the detector performance for this decay mode. In addition, the analysis was performed at luminosities of $3 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ and $6 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$. All results are shown in Table 4-6 and the invariant mass plots for the $6 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ case are in Figs. 4-10(d)–(f).

Table 4-6
Summary of efficiencies and number of fake tracks for $H^0 \rightarrow e^+e^-\mu^+\mu^-$ events for various configurations.

Luminosity	Fakes per event with $p_t > 5$ GeV/c	Track efficiency $p_t > 10$ GeV/c	Electron E/p efficiency $0.7 < E/p < 1.4$	M_Z cut efficiency		Higgs reconstruction efficiency
				e	μ	
1×10^{33}	0.03 ± 0.01	0.991	0.96 ± 0.01	0.99 ± 0.01	0.99 ± 0.01	0.84 ± 0.04
3×10^{33}	0.04 ± 0.02	0.989	0.96 ± 0.01	1.00 ± 0.01	0.97 ± 0.01	0.83 ± 0.04
6×10^{33}	0.18 ± 0.03	0.972	0.93 ± 0.01	1.00 ± 0.01	0.93 ± 0.02	0.75 ± 0.04

It is seen that the tracking detector should allow excellent performance for Higgs reconstruction well beyond the design luminosity. This is the single most important design goal for the detector. The track reconstruction was also performed without using information from silicon barrel layers 1 and 2 (at $r = 9$ and 12 cm respectively) to check that the design is robust against some inefficiency. The efficiency reduction for this case is about 10%, even at six times design luminosity. Note that occupancy in the STS is not a problem at six times design luminosity and thus does not impair high luminosity performance. The efficiencies for finding segments in the individual OTD superlayers are indicated in Fig. 4-11 for Higgs

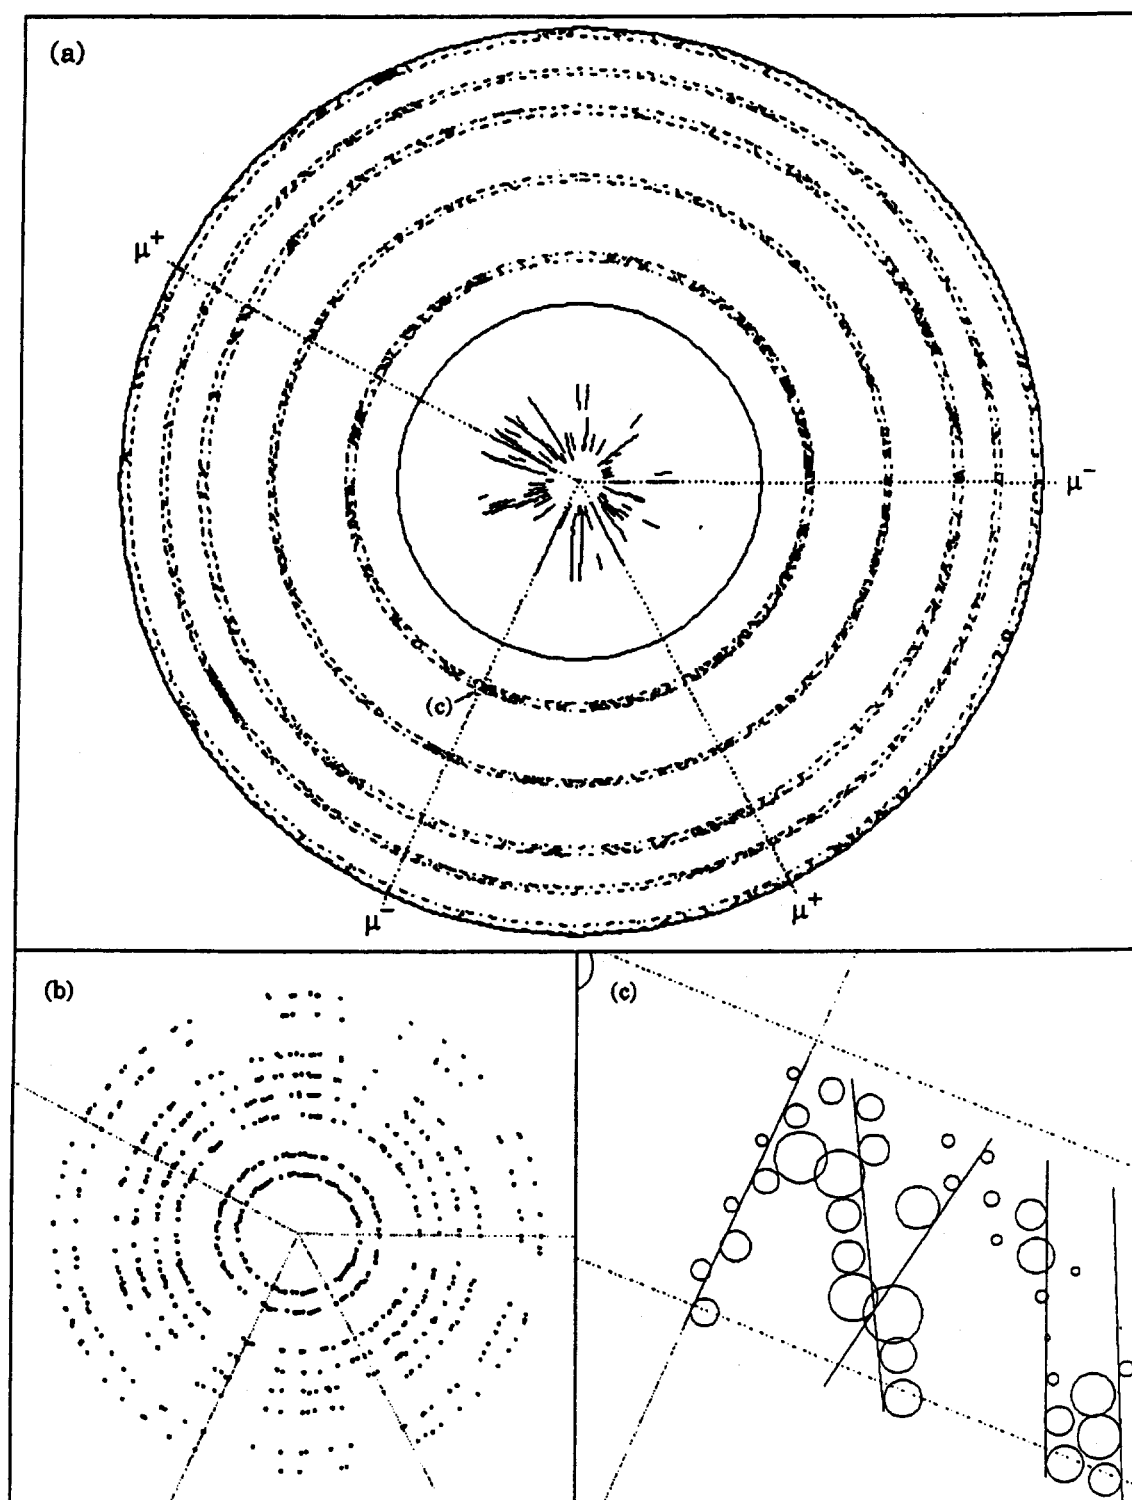


FIG. 4-8. a) The barrel region of the central tracker for a leptonic Higgs event at design luminosity, showing the reconstructed muon tracks from the Higgs decay, the raw data for the outer tracks and reconstructed track segments from the silicon tracker. b) Close-up of the barrel region of the silicon tracker, showing reconstructed space points. c) Close-up of the inner straw superlayer corresponding to the marked region in a). Straw tube hits are represented by circles with radius proportional to drift time. Short lines represent locally reconstructed track segments, and the long line is a fully reconstructed high transverse momentum track.




FIG. 4-9. Efficiency of single tracks with $p_t > 1$ GeV/c in Higgs events as a function of pseudorapidity.

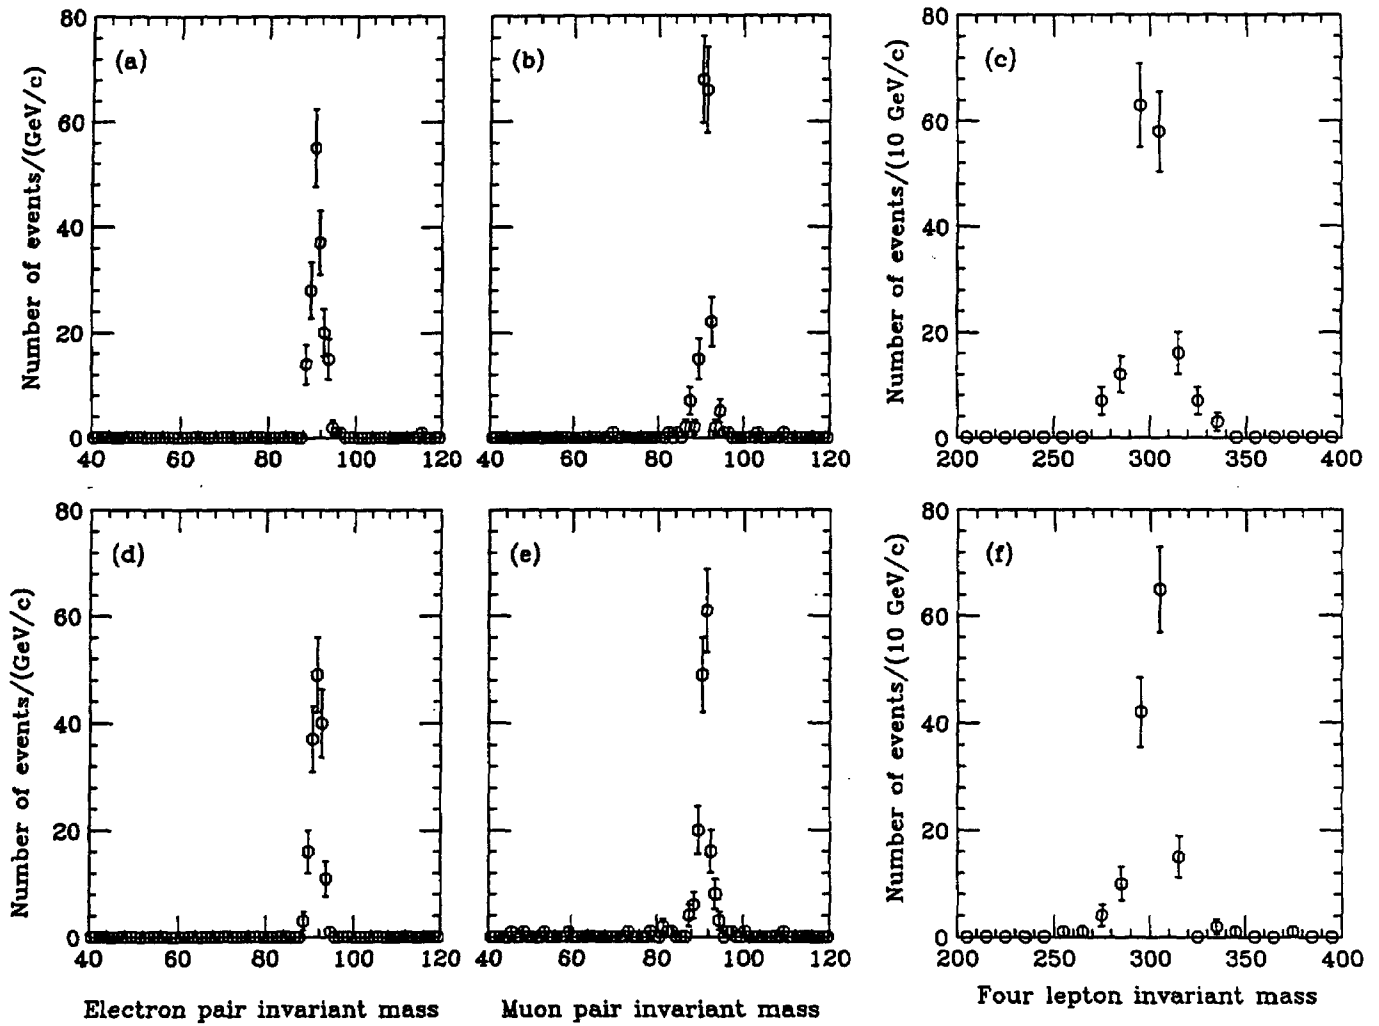


FIG. 4-10. (a) Dilepton invariant mass distribution for electrons that have passed the $P_{\text{gen}}/P_{\text{fit}}$ cut. (b) Dilepton invariant mass distribution for muons. (c) Reconstructed four-lepton invariant mass distribution for a 300 GeV Higgs. (d)-(f) similarly for $6 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. All masses are in GeV.

events (to 4 muons in this case) with no background and for background levels corresponding to 1-6 times design luminosity. All tracks from the Higgs-producing interaction having $p_t > 2$ GeV/c are considered "interesting" tracks for this evaluation. Tracks with lower p_t , down to 1 GeV/c, have also been examined at design luminosity; the efficiency is similar, but the right/left ambiguity resolution is less reliable in the outer superlayer because 1 GeV/c tracks cross at large angles relative to the radial direction. The high efficiency in the outer layers over the full luminosity range guarantees excellent momentum resolution when the combined OTD and STS measurements are used.




FIG. 4-11. Efficiency for finding segments of single tracks with $p_t > 2$ GeV/c for each superlayer of the OTD for $H \rightarrow 4\mu$ with background levels corresponding to the luminosities indicated.




FIG. 4-12. b tagging efficiency vs. jet p_t for the combined system using the method described in the text.

Capability for b -jet tagging with secondary vertices

The capability to identify b -quark jets is expected to be an important tool for studying the $t\bar{t}$ system, by reducing the non- $t\bar{t}$ background substantially, and also reducing combinatoric background for measuring the mass of the t quark. The efficiency for tagging b jets using secondary vertices has been investigated with the GEANT simulation and full event reconstruction [7]. A simple algorithm is used, which is to count for each jet the number of reconstructed tracks with significant impact parameter. Figure 4-13 shows a simulated $t\bar{t}$ event where both W bosons from the t decays subsequently decay leptonically.

ISAJET is used to generate b and "generic" jets in separate jobs, and tracks are reconstructed with the combined silicon and straw system. Tracks with measured $p_t > 0.8$ GeV/c within the cone $\Delta R < 0.5$ to the jet axis are associated to each jet. With the algorithm used, the reconstruction efficiency for these tracks ranges from 95% (jet $p_t = 50$ GeV/c) to 80% (jet $p_t = 500$ GeV/c). Tracks are further required to have $\chi^2/DOF < 5$ and $b_0 < 1$ cm (in order to discriminate against reconstruction errors and secondary interactions), and the impact parameter b_0 is required to have the correct sign to be consistent with originating from the decay of a particle moving in the jet direction.

Figure 4-12 shows the tagging efficiency for b jets compared to that for a mix of jets from gluons and u , d , s and c quarks, where a tag requires 3 or more tracks with impact parameter greater than 3σ . For b jets with 40 GeV/c $< p_t < 200$ GeV/c, the efficiency is approximately 30%, and the rejection of the non- b jets is of order 100:1. Using these b tagging results from the full GEANT simulation, the efficiency for tagging one or both b jets from $t\bar{t}$ decay is estimated to be about 50%, assuming that only secondary vertices are used for the tag.




FIG. 4-13. GEANT simulation of an Isajet $t\bar{t} \rightarrow e\mu\nu\nu b\bar{b}$ event with $M_{\text{top}} = 150 \text{ GeV}/c^2$, showing the beam pipe, inner tracking gas enclosure, and first barrel silicon layer at 80% of actual size. Tracks reconstructed using data from the silicon and straw systems with $p_t > 2 \text{ GeV}/c$ and $|\eta| < 2.5$ are plotted. Charged leptons and b jets from the $t\bar{t}$ decay are indicated with their generated momenta; the muon is actually inside one of the b jets. Arrows indicate the true decay vertices of the b and c mesons within the b jets. The secondary vertices are visible by eye from the reconstructed tracks.

The SDC design will allow b tagging over a very broad range in jet momentum. This is relevant not only to the $t\bar{t}$ studies, but could be of interest if new physics is discovered which populates preferentially the heavy quark sector.

4.2.2. Material in the tracking volume and electron identification

The SDC tracking system differs in at least two significant ways from previous tracking systems in collider experiments: the tracking technologies that are matched to SSC physics have intrinsically more mass, and the necessity of providing tracking over the rapidity interval $|\eta| < 2.5$ means that much of the mechanical support material for the tracking system is inside the tracker. As a result, the SDC tracking system will have more material than previous collider trackers. In this section we present the main results of the SDC study on material in the tracking volume [8] as they relate to the design and optimization of the tracking system. Although the designs that were considered in Ref. 8 are not identical to the one in this report, they are close enough that the conclusions can provide guidance for our design.

The material in the tracking system impacts the SDC trigger, and the reconstruction and identification of electrons and photons as follows:

1. *Triggering:* Electron bremsstrahlung and photon conversions will occur in the material in the SDC tracking system. Electron bremsstrahlung will put some high- p_t electrons below the Level 1 tracking trigger threshold resulting in a smearing of the p_t threshold. Photon conversions will result in a number of apparent high- p_t electrons that will pass the Level 1 trigger, and may result in a loss of prompt high- p_t photons at the trigger level.
2. *Momentum resolution:* Bremsstrahlung in the material will lead to degraded momentum resolution for electrons.
3. *EM calorimetry:* Electron showers are identified by requiring that Had/EM is small, and that transverse shower profiles are consistent with an EM shower. Bremsstrahlung in the tracking system can result in several separated energy deposits in the EM calorimeter rather than one from the electron shower. This results in altered transverse shower profiles, as measured in EM towers and in the strip detectors at shower maximum. As a result, electrons may be identified with reduced efficiency.
4. *e/γ ID:* e/γ separation is achieved by identifying a charged track pointing at the EM shower and separation of electrons from overlapping photons, and charged hadrons is achieved by track-shower matching using the strip detectors at shower maximum, and by requiring a selection on E/p near 1. Degradation of the electron momentum resolution and the EM calorimeter cluster position resolution will result in reduced efficiency for e/γ identification.

The conclusions of Ref. 8 are that the most significant impact of the SDC tracking material is through e bremsstrahlung. As noted above, bremsstrahlung smears the e trigger threshold(s) and reduces the efficiency for identifying electrons cleanly. The effect of the material is worst for the tracking elements closest to the beam. Thus it is recommended that the material in the silicon tracker (including the beampipe and supports) be not more than 6% X_0 at $\eta = 0$, and $\leq 7\%$ when averaged over rapidity. The use of the silicon tracking information in the Level 2 trigger should allow significant suppression of the electron rate from conversion pairs. This is not needed at Level 1.

The remaining tracking material should also be minimized, since material before the Level 1 tracking trigger elements results in a "softening" of the high- p_t electron trigger. The goal for the tracking system is to have $\lesssim 10\%$ of X_0 at $\eta = 0$ and $\lesssim 15\%$ averaged over η , excluding any layers directly in front of the calorimeter. With this level of material before the Level 1 tracking trigger elements, the rates for isolated electrons from the conversion of QCD prompt photons are tolerable relative to the physics backgrounds: $W \rightarrow e$, $t \rightarrow e$ etc.

Unlike previous collider tracking systems, the SDC inner silicon tracking system is by itself capable of full track reconstruction and is sufficient to make a determination of the track momenta and directions. This is useful for the optimum identification of electrons in SDC since the use of the momentum measured in the silicon provides the most efficient E/p cut [9]. For example, a cut with $E/p < 1.4$ allows typically a 95% efficiency for keeping electrons.

The material in the tracking volume has been calculated using a GEANT based program. For the STS, this program includes: silicon detectors, support cylinders, beampipe and enclosure walls, electronics

packages, cooling structures, cables, and the gas volume. Minimizing the amount of material requires a careful choice of support structures.

The OTD and ITD support systems similarly must provide a rigid and precisely aligned structure with the smallest possible amount of material. For the OTD, this is accomplished by using a carbon fiber composite and foam laminate for the module shells and for the cylinders that support the modules. The cylinders are supported by a carbon composite space frame. The amount of material in the outer tracking system for $\approx 90^\circ$ incident tracks is 3.5% of a radiation length, including all supports, but not including the last superlayer. This gives an amount of material $3.5\%/\sin\theta$, where θ is the polar angle.

The ITD structures, although less well defined at present, are expected to have a thickness at normal incidence of 7% of a radiation length on the average. This gives for this device a material distribution which is $7\%/\cos\theta$.

The number of radiation lengths traversed by a particle through the entire tracking system is shown as a function of rapidity in Fig. 4-14. Included are the beam pipe, the silicon strips and their associated readout electronics, supports, and utilities, the straws and their support cylinders, gas, wires, end structures, readout boards and cables, as well as the gas microstrip disks and their support. We discount the outermost straw superlayer, the last disk, and the readout boards of the outer three cylinders, since particles leave the system and enter the calorimeters or coil immediately after traversing these. The dotted curve below the peak around $|\eta| = 1.9$ indicates the effect of omitting the blocking capacitors for the straw tubes. The contribution of the barium titanate capacitors currently in hand (and modelled in the calculation shown) is large enough to warrant significant R&D effort toward a lower mass solution. The material averaged over η from 0 to 2.5 is 13.5% of a radiation length.




FIG. 4-14. Number of radiation lengths vs. rapidity for the full tracking system (solid curve) and silicon alone (dashed curve).

4.2.3. First level trigger

The SDC trigger design calls for each subsystem to process the raw data from the detector signals for each crossing into signatures for specific trigger "objects." These objects are composed from data packets that are sent via fiber optic cables to the Level 1 trigger decision logic located in the surface building above the detector. The selections that are made on the raw detector outputs and the information that is included in the data packets represent specific choices [10,11] designed to select interesting physics while rejecting the backgrounds. For tracking, the selection imposes a lower limit on the track p_t and the information includes the location and momentum threshold crossed. The goal of the Level 1 tracking trigger is the rapid and

efficient generation of trigger information and the efficient encoding and transmission of the information to the trigger decision logic. For such track related objects to be manageable in number, a lower limit must be imposed on their momenta, or an association must be made with some other detector element that provides this limiting restriction. The important aspects of the track trigger are 1) efficiency, 2) transverse momentum resolution, 3) positional resolution (granularity), 4) false rate, and 5) time resolution. The timing requirement for the generation of the Level 1 trigger is dealt with in detailed reports by the Trigger Group [12,13].

The OTD superlayer trigger

A high momentum track in a constant solenoidal field has a trajectory which, in cylindrical coordinates, is approximately represented by

$$\phi = \phi_0 + K\tau + b/\tau \quad z = r \tan \lambda + z_0,$$

where the track parameters are ϕ_0 , K , b , $\tan \lambda$, z_0 . For a track emanating from the origin, the impact parameter $b = 0$. The parameter $K = 1/(2 \text{ radius of curvature})$ measures the transverse momentum and can be locally calculated in a superlayer through the relation $d\phi/d\tau = K$. This is the basis for the Level 1 trigger using a straw tube superlayer. For $B = 2 \text{ T}$, the goal for a Level 1 trigger from Table 4-1 corresponds to measuring K with an error $\sigma_K = 3 \times 10^{-3} \text{ m}^{-1}$.

The proposed trigger design is based on 1/2 cell staggered axial superlayers of 8 straws which are radially aligned and closely packed [14]. A trigger cell is 1/2 straw wide and is formed from 2 groups of 4 radially aligned straws. In this design the arrival times of signals from the consecutive layers of straws within the superlayer are processed electronically to select only tracks above a preset lower momentum limit, called a stiff track. The timing relationship between the arrival times of the 1/2 cell staggered straws provides the necessary information to resolve the left-right ambiguity and to eliminate the timing uncertainty due to the drift intervals of the individual straws. What remains is the timing uncertainty due to the particle flight path and signal transmission time along the straw. The combination of these two uncertainties is about 8 ns and falls within a single crossing time of 16 ns. Therefore, individual crossings can be specified for each high momentum track sensed. The timing resolution and lever arm available within one superlayer of straws would imply a momentum resolution of 20% at 10 GeV/c if only one layer were used.

The processing of signals from the straws for triggering will be implemented using signals from the straw discriminator outputs. The circuit that processes the arrival times is composed of numerous digital mean timers [15] and has been implemented in CMOS and verified to function as needed. Output from the trigger chips, stiff-track signals, will pass to the global trigger logic through an information path independent of the data acquisition path.

The behavior of the proposed trigger has been simulated with GEANT and the SDC geometry for H decay to WW . The simulation includes minimum bias events from 4 previous and 2 following crossings. Real tracks result in an 8 ns wide peak located at precisely the point expected for the circuit configuration simulated. The background of minimum bias events and incorrect combinations of wire hits is characterized as a false rate per trigger cell per crossing.

The trigger circuit threshold in transverse momentum is programmable up to a limit imposed by the electronic range accepted by the design. While the circuit design chosen can accommodate thresholds down to 5 GeV/c, it is expected that a threshold of 10 GeV/c will be used. Simulation of the circuit performance was done with a 10 GeV/c transverse momentum limit. The trigger selection shows some leakage due to modest momentum tracks within jets that have accompanying hits from overlapping soft tracks. The false triggering rate from minimum bias events is, according to the simulation, 1 for every 4 crossings at $10^{33} \text{ cm}^{-2}\text{s}^{-1}$. Thus for each of the 128 trigger cells at the outer superlayer, 1/512 of the crossings have a trigger from the overlap of minimum bias hits. The rate for real tracks over 10 GeV/c from minimum bias events is considerably below this number.

Use of several superlayers in the Level 1 trigger

Since the superlayer trigger objects all arrive at the central Level 1 decision logic, use of several superlayers is possible either as an original plan or as an upgrade as the luminosity is increased. The false rate can be suppressed and/or the efficiency can be enhanced by the use of more superlayers. For this reason simulations that make use of the primitives from the three axial superlayers have been performed. A trigger based on a two-out-of-three combination allows both a significant improvement in the rejection of fakes and a higher efficiency when compared to a single superlayer trigger. This higher efficiency is necessary to meet the goals for the tracking trigger. Figure 4-15 shows the threshold curve for the two-out-of-three trigger which is the baseline design. This trigger algorithm has a dramatically reduced false trigger rate when compared to one-per-four-crossings from the single superlayer trigger. In 500 minimum bias events simulated there were no triggers. This result was derived from the same Higgs event sample discussed above by setting the crossing gate to the crossing prior to that which contains the Higgs event. This further illustrates the precision with which the crossing can be determined by the mean timers.




FIG. 4-15. Threshold curve for the two-out-of-three superlayer OTD Level 1 trigger.

ITD Level 1 trigger

For the ITD the coordinate measured directly is ϕ vs. z . This allows a local measurement $d\phi/dz = K/\tan\lambda$. Binning in η provides a rough determination of $\tan\lambda$, so that the measurements can be used directly to cut on the transverse momentum. Since $\tan\lambda$ can be nearly 10, the ITD needs 10 times the tracking length in z as compared to the straw superlayer radial extent to achieve a comparable error on K , for comparable errors in ϕ measurement. In fact, since the error on ϕ is proportional to the position error σ divided by r , an ITD superlayer typically measures ϕ less well than a straw superlayer. Adding errors from the determination of $\tan\lambda$ and the nonuniformity in the B field results in the necessity of a large separation in z for the elements making up the ITD Level 1 trigger. The major problem facing the ITD trigger is therefore the collection of a large amount of data from spatially separated units to a central location in a sufficiently short time to be useful at Level 1. Unlike the case of individual straws, the ITD signal for each measuring element will occur within the time bucket of the event since both the drift time and electrode propagation time are very small.

The SDC has about a 1 m separation for the elements used in the ITD trigger. A minimum amount of information for such a trigger could be provided by two superlayers, using $(\phi_1 - \phi_2)/(z_1 - z_2)$ as the quantity in the trigger. An additional superlayer to provide confirmation is included since a two-layer scheme is very susceptible to background from random combinations of hits.

The ITD trigger assumes a division of the η range into 4 bins, with projective geometry for the detector elements. Each bin in η provides an average value of $\tan \lambda$ to be used to relate $d\phi/dz$ to K . The turn-on curve for the trigger in the different η bins is shown in Fig. 4-16. The background is calculated to be very small.




FIG. 4-16. Trigger threshold curve for the four separate η bins of the ITD trigger.

4.2.4. Second level trigger

The Level 2 tracking trigger has as its goal the tightening of momentum errors for electron and muon candidates, the elimination of fakes by the inclusion of more information, and the elimination of a majority of the conversions contributing to the electron trigger. To accomplish these goals, the second level trigger uses all tracking devices, including the silicon. The momentum cuts are still based on the change in azimuthal coordinate with either the radius or the z coordinate in η bins. The linking of data from several superlayers allows the elimination of nearly all fake tracks. In addition, the presence of a track throughout most of the silicon detector should allow rejection of about 70% of the conversion electrons, with the remainder dominated by internal conversions. To reject most of the rest of the conversions would require the use of analog information, an upgrade option, in the inner silicon layers with a pulse height cut to keep single hits and reject double hits from pairs.

4.2.5. Alignment requirements

The alignment requirements for the tracking system have been described in detail in Ref. 16 and are summarized here. We describe the geometry by a misalignment vector at every point in the detector which describes the displacement at the point from an ideal data base of positions established through some calibration procedure. This vector can be resolved into components in cylindrical coordinates, i.e., circumferential, radial, and longitudinal displacements. The alignment requirements are specified by the allowed values for these components, both locally and correlated over the whole detector. Since the transverse momentum is measured in a solenoidal field by the change in circumferential position vs. radius, the requirements are significantly more stringent on this coordinate.

The local alignment requirements for the coordinates directly measured can be specified by comparison to the precision of each measurement. A reasonable goal is a misalignment error $\lesssim 40\%$ of the measuring error. For the nonmeasured coordinate (e.g., the assumed radius of an intercepted cylindrical layer) the misalignment results in a displacement of the measured coordinate dependent on the track parameters.

Again, limiting this error relative to the measuring and multiple scattering errors will lead to a specification for the misalignment in this coordinate. Finally, by specifying the dependence of the momentum on the full set of measurements, the effect of correlated misalignments can be quantified.

Based on such an analysis, the specifications in Table 4-7 (always RMS values) have been set for uncorrelated and several correlated misalignments.

Table 4-7
Alignment requirements.

Maximum local misalignment:		
Silicon:	5 μm	Circumferential
	250 μm	z (barrels) or r (disks)
	80 μm	r (barrels) or z (disks)
Straw superlayer:	35 μm	Circumferential
	250 μm	in z
	1200 μm	in r
Gas microstrips:	40 μm	Circumferential
	1000 μm	in z
	250 μm	in r
Global:		
Azimuthal rotational alignment of silicon vs.		
straws or gas microstrips		10^{-5} radians
Common centering silicon vs. straws		15 μm
Common centering silicon vs. gas microstrips		40 μm
Centering of tracker on beam		500 μm

To achieve the numbers in this table we assume that the silicon placement accuracy is 25 μm in the circumferential direction with a calibration procedure using surveying and x rays providing the location to 5 μm . The mechanical system is then required to be stable over time to 5 μm .

For the straw system, individual straw placement errors within modules are expected to be small and the most stringent requirement is on the positioning of each module as a unit. The various tolerances making up the overall module placement error are described in section 4.6.5. Surveying and x ray measurements will be used to verify straw and module positions and establish an initial data base of locations. Stability over time of the mechanical structure will be important for achieving the alignment goals.

The global alignment goals will be very difficult to achieve using surveying techniques alone. We plan to also use high p_t tracks to establish a data base of the relative detector positions. It is critical, however, that the mechanical structure and positioning are stable over time so that the global alignment can be specified in terms of a small number of parameters such as the locations of the centers and relative rotations of the detector subsystems.

A final alignment issue arises in the case of triggering where we cannot input alignment constants, as in the case of offline reconstruction. The assumption of the trigger is that tracks originate from the origin. Thus a beam offset is equivalent to an impact parameter for a track, which yields locally a spurious apparent curvature. Using the equation for a stiff track

$$\phi = \phi_0 + Kr + b/r,$$

gives

$$\frac{d\phi}{dr} = K - \frac{b}{r^2},$$

where b now is the beam offset orthogonal to the track direction at the origin. Thus we need to keep

$$\frac{b}{r^2} < \sigma_K.$$

For the Level 1 trigger, the goal for σ_K is $3 \times 10^{-3} \text{ m}^{-1}$; while for the Level 2 it is a factor of two better. Using typical radii for the systems (25 cm for the silicon, 50 cm for the gas microstrips and 150 cm for the straws) we get the following requirements:

$$b < 7 \text{ mm} \quad \text{for straws at Level 1}$$

$$b < 800 \text{ } \mu\text{m} \quad \text{for the gas microstrips at Level 1}$$

$$b < 100 \text{ } \mu\text{m} \quad \text{for the silicon alone at Level 2.}$$

These last two numbers are quite stringent. A significant reduction to the sensitivity on b for the silicon alone can be achieved by using the full tracker in the Level 2 trigger. In the barrel region, a more detailed analysis for the full set of tracking detectors results in the need to keep $b < 500 \text{ } \mu\text{m}$.

4.2.6. Possible upgrades

The SDC baseline tracking design is planned for a decade of operation at the SSC. The SSC program is, however, expected to be very long lived and upgrades of the present system emphasizing vertexing, pattern recognition, momentum resolution or very high luminosity are possible and may be desirable as we learn about physics at the TeV scale. The tracking detector is probably the most easily upgradable subsystem of the SDC.

Some of the possible options, with their strengths, are as follows:

- a) Improvements in vertexing through the evolution of pixel detectors.
- b) Improved momentum resolution and luminosity capability by deploying silicon strip detectors at larger radii.
- c) Improved luminosity performance in the barrel region through the use of scintillating fibers or gas microstrips—particularly for the innermost layers.

A scintillating fiber detector, which remains a possibility for our final baseline design, is described in detail in section 4.8.

The devices mentioned above are those closest to being realizable at present. It is possible that even more options will exist in another decade.

4.3. Radiation environment

4.3.1. Fluence of particles from interactions

The tracking detectors experience a large fluence from particles produced directly from the beam-beam interactions and albedo from subsequent particle interactions in the calorimeters. With a polyethylene endcap liner to suppresses albedo neutrons, the fluence at radii less than 50 cm is dominated by charged particles from beam-beam interactions. For larger radii the fluence is dominated by neutrons, but both the straw detectors and gas microstrips are relatively insensitive to these. Thus here again the charged particles are the major concern. For the straw tubes these create occupancy, and for the gas microstrips they can create problems by charging the detector substrate. For the silicon, the detector performance is eventually degraded because of displacement damage in the crystal.

For a uniform rapidity distribution of produced particles, with no magnetic field, the charged particle fluence varies as $1/r^2$, where r is the distance from the beamline. With a magnetic field, charged particles are trapped, increasing the particle fluence. Particles from photon conversions further increase the occupancy but produce little damage in the silicon detectors, since electrons create very little displacement damage to the crystal lattice. With the coil radius of 1.7 m and a B field of 2 T, all primary charged particles with transverse momenta below 510 MeV/c will be trapped. The resulting fluence increase depends on radius, does not vary greatly with z , and is typically a factor of 1.5 to 2. The fluence of charged hadrons expected for a year of running at the standard SSC luminosity is shown in Fig. 4-17, as a function of the radial distance, for the region occupied by the silicon.




FIG. 4-17. Charged hadron fluence vs. radius for one year at $L = 10^{33} \text{ cm}^{-2}\text{s}^{-1}$.

4.3.2. Albedo neutrons

Calorimeter albedo is the principal source of neutrons (typically 1 MeV) inside the SDC tracking volume. Following the prescription of Ref. 17, we have simulated the spatial distribution of the neutrons from the SDC calorimeter considering a cylindrical cavity of half length 420 cm, radius 205 cm, and beam hole of radius 42 cm at $z = 420$ cm. Details of the simulation are given in Ref. 18.

The average number of reflections from the calorimeter face was assumed to be equal to one for an unlined calorimeter. For calculations of yearly fluences, the luminosity assumed was $10^{33} \text{ cm}^{-2}\text{s}^{-1}$ for 10^7 s. Figure 4-18 shows the fluences for radii $r = 0, 50, 100$, and 170 cm, respectively. The increase for large z , close to the beam hole, is obvious.

To reduce the neutron fluence, particularly at large z , we will line the faces of the endcaps with 12 cm polyethylene liners, giving a reduction of a factor of 5 in the direct albedo and a total suppression of the reflected albedo from the lined surface. In our simulation, shown in Fig. 4-19 for a 10 cm liner, we get a reduction of the neutron fluence of about 2 in the low- z region and up to 5 close to the endcaps. Moreover in this case the fluences become extremely uniform over the whole cavity. Typical yearly fluences are:

	No liner	Endcaps lined
all $r, z < 250$ cm	6×10^{11}	3×10^{11}
$r = 100$ cm, $z = 400$ cm	2×10^{12}	4×10^{11}

It has also been suggested that the liner include a thin layer of boron to absorb neutrons below 100 keV. We intend to investigate this idea in the future.




FIG. 4-18. Yearly neutron fluences in the SDC tracking volume, at $L = 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.




FIG. 4-19. Yearly neutron fluences in the SDC tracking volume with endcap calorimeters lined with 10 cm polyethylene at $L = 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.

4.3.3. Backsplash from magnet coil

Using the GEANT-based SDC tracking simulation the occupancy of the straw tube superlayer of the OTD nearest to the magnet coil was studied. Occupancy is defined as the fraction of wires in a given layer recording a hit when the system is read out. Since low-energy particles can loop through the system for extended periods of time, simulated interactions were generated over the range of -4 to $+2$ beam crossings relative to the time of the bucket studied. The hits from these interactions were recorded using a 50 ns gate. The simulated interactions were minimum bias events since the issue of interest was overall occupancy. The simulation allows one to switch the mass of the solenoidal magnet off while leaving the magnetic field within the tracking volume unaffected. Two samples of about 100 simulated event read outs were studied: one with and one without the magnet mass. The results of this study showed that the magnet coil contributed a small (possibly statistically insignificant) rise in occupancy from 0.026 to 0.029 for the straws in the outer superlayer. This indicates that the effect of the magnet material is not great.

4.3.4. Detector lifetime

The STS sits in the most harsh radiation environment of all of the tracking devices. Therefore, the study of the effect of radiation on silicon detectors has been a very important area of R&D over the last two years.

Table 4-8 shows the expected yearly fluences from charged hadrons at design luminosity for various radii of the silicon tracker. We ignore the effect of neutrons since their contribution ($3 \times 10^{11} / \text{cm}^2$) is expected to be small, as discussed earlier.

The detector lifetime is expected to be about 10^{14} particles/ cm^2 for 1 MeV neutrons or 800 MeV protons as shown by our measurements [3]. For the particles at the SSC, mainly relativistic pions, the damage is on average a factor of two smaller. We take this factor as a safety margin on our calculations. We see from the table above, that the STS has 7 barrel layers at all rapidities with lifetimes ≥ 10 years at Standard Luminosity. The inner most barrel layer is the only one with an expected lifetime less than 10 years. All forward disks have lifetimes ≥ 17 years.

An additional concern is the effect of radiation on the STS support structure and its impact on the detector alignment. We have tested graphite epoxy corrugated panels, which form the primary structure, to verify their dimensional stability under radiation.

Table 4-8
Yearly fluence vs. radius.

Radius (cm)	Fluence particles/cm ²	Number of years to reach 10 ¹⁴ /cm ²
10	11×10^{12}	9
15	5.9×10^{12}	17
20	3.3×10^{12}	30
25	2.2×10^{12}	45
45	0.95×10^{12}	105

The OTD is located just inside the magnet coil and no closer than 80 cm from the interaction point where the total dose of radiation is expected to be about 100 Krad and the total charge deposited on a chamber sense wire is expected to be about 0.2 C/cm over a ten year time span at the design luminosity. The materials of the straw tube tracking system, mylar, Kapton, epoxy, carbon fibers, etc., are known to be tolerant to radiation levels of this magnitude. The cabling can also be made with radiation tolerant cables. Excluding the electronics, which is discussed elsewhere, the survivability of the OTD in the radiation flux is therefore not limited by its materials.

The limits to the survivability of the straw tube tracking system come from the action of electric currents on the anodes and the cathodes of the straw tubes themselves during operation. A gas consisting mainly of tetrafluoromethane (CF₄) with the addition of isobutane at the level of 20% will be used. This gas is known to cause very little loss of gain at the anode even up to levels of deposited charge far beyond the expected amount at the SSC. Studies [19,20] have shown effectively no anode damage up to at least 2 C/cm. This corresponds to more than 100 yrs at the design luminosity for the inner layer. The limitation of the lifetime due to effects at the anode is therefore insignificant.

The cathode is potentially a more significant site for radiation induced damage than the anode in CF₄/isobutane. The cathode will, however, be made of copper coated Kapton rather than an aluminum coated film. Copper coated tubes do not show the same tendency for sustained discharges to occur that has been observed in aluminum coated tubes at the level of about 0.1 C/cm. The radiation survivability of small tube samples has been tested to 0.3 C/cm, which is what would be expected in the inner(outer) superlayer from 15(60) years of operation at design luminosity. At the end of this period of irradiation no indication of any difficulty was observed. The voltage on the tube could still be raised 50% above the operating voltage with an intense flux of particles passing through the tube generating about 10 μ A/cm. It appears that the lifetime for the inner layer is actually far in excess of 15 years. More irradiations are planned to increase the measured levels of survivability.

The limitation imposed by the gradual loss of cathode material due to ion bombardment is on the order of 0.001 μ m per year for the inner layer at design luminosity compared to a typical coating thickness of 0.15 μ m. An additional thickness can be applied to provide a 50 year period before the cathode resistance increases above its design value.

The ITD sits in a region where the radiation dose is expected to be typically 100 Krad over a 10 year period at the standard luminosity. The ITD ionization dose may affect the detector substrate electrical and mechanical properties. This is the most important R&D area for gas microstrip detectors. The issues and plan for R&D are outlined in detail in section 4.7. The goal is to have a stable substrate with ohmic resistance of about 10¹³ Ω per square. Once this issue is resolved the gas microstrips are expected to be very radiation hard.

4.3.5. Electronics lifetime

All tracking systems are planning to use bipolar analog front-end electronics and CMOS digital data storage chips. The choice of bipolar is based on its excellent noise performance and low power at high speeds. CMOS is necessary where very high density circuitry is needed.

Small-feature-size fast bipolar transistors have been shown to be intrinsically radiation hard. Data on the gain stability of AT&T npn transistors (run at the low currents needed for low power operation) for both ionizing radiation and neutrons are shown in Fig. 4-20(a) and (b). Data also exist on the radiation hardness of other fast bipolar processes: AT&T [21], NTT [22], AT&T, VTC, and Harris [23], Tektronix [24].




FIG. 4-20. (a) Gain stability of AT&T npn transistor after ionizing dosage; (b) gain stability of AT&T npn transistor after neutron irradiation.

In addition to the AT&T process, the Tektronix process is attractive because of the very small feature size. The irradiation data show that the current gain of npn transistors with small base spreading resistance ($r_{bb} = 60 \Omega$) like the (Tektronix SHPi 8x) starts out at $\beta = 88$ at $I_c = 100 \mu A$ and still yields $\beta = 41$ after irradiation of 1.1×10^{14} 650 MeV protons /cm². For this process JFET's, resistors and diodes are extremely radiation resistant.

Data also exist for pnp transistors, which are desirable as active loads (an alternative would be JFET's). The AT&T vertical pnp transistors are radiation resistant with current gain $\beta \approx 20$, which is certainly adequate, after either 2 Mrad or 1×10^{14} neutrons/cm². The Tektronix lateral pnp transistors show degradation down to $\beta \approx 5$ after the same fluences; this still might be good enough to be used for active loads.

We now have good data to predict radiation effects and incorporate them in the circuit design. We have not observed a deterioration of the very good noise performance of bipolars over a wide band of frequency. The ultimate choice of the specific bipolar foundry will be a question not only of performance but also of transistor size, yield and price.

We have extensively tested the performance of radiation-hard CMOS for analog applications (noise, transconductance [25]) and have found that the bipolar option is superior for the fast shaping times needed. For digital applications, the radiation-hard 1.2 μm CMOS process from UTMIC looks very promising [26]

with threshold shifts of about 150 mV after irradiation with about 10^{14} protons /cm² and smaller shifts in logic gates. The digital part of the silicon front-end system (digital data storage) will be built in radiation-hard CMOS.

4.4. Design of detector elements

All of the devices making up the tracker are constructed from modules which can be separately assembled, powered, and tested. These modules then have to be assembled into a highly stable mechanical structure with integrated cabling, cooling, and electronics. We describe the individual module designs in this section.

4.4.1. Silicon detector and module

The detector subassemblies listed in Table 4-2 have been chosen to fit with the choice of a 12 cm long detector unit. The central region barrel detector can be made of 5 detector units, each 12 cm long. The forward disks are generally made of either one or two 12 cm units and require the fabrication of eight detector types. For the disks the electronics are assumed to be at the outer periphery, where the pitch is approximately 50 μ m. The position accuracy thus improves at radii smaller than the maximum for such a 12 cm disk module.

Figures 4-21(a) and (b) show the configuration of barrel detectors and the largest disk detector as fabricated using a 4 inch silicon wafer. The detectors have been arranged to lie within a 9 cm circle to optimize yield in fabrication. The barrel detectors are made two to a wafer, the disks one to a wafer. The barrel detectors have 5 128-channel chips per side for the readout, the disks 10 per side. Table 4-9 lists the number of 6 cm long detectors per layer, the number of readout channels per layer, the number of modules in ϕ per layer, and the detector overlap (to minimize dead areas) per layer.

We summarize in Table 4-10 the specifications for the individual SDC double-sided silicon detectors, for the barrel geometry. The disk detectors are rather similar except that they are about twice as wide. Table 4-11 lists the dimensions of the active area of the various disk modules and the pitch on the inside and outside.

Figure 4-22 shows the pattern of structures for each side of a detector. The detectors are AC coupled, so that the front-end electronics do not have to handle large leakage currents. The implants and readout strips are narrow to minimize capacitance, and the *n*-side isolation is provided by *p*-implants which are chosen to be wide, in order to minimize capacitance. These structures have been shown by our tests to maximize radiation hardness. Figure 4-23 shows the steps in constructing a detector module from the individual detectors. This involves gluing detector units together into a mechanical unit, stiffened by edge strips. The readout units are then bonded together, with electronics and cables at the ends.

Module layout

Integration of the detectors, readout ICs and connecting cables is a critical area where definitive results are not yet available. Although a substantial body of semi-quantitative work has established guidelines for the electronic and mechanical design, very careful and detailed measurements are necessary to verify crucial aspects of the concept. The following discussion outlines our current design approach, which has resulted from several iterations of interdependent electronic and mechanical designs. We are now focussing on a set of critical measurements to verify the design concept. This work is a top priority. Critical issues are cross-talk from the readout chips to the detector and reduction of added material to a small fraction of the detector silicon.

The baseline design utilizes aluminum/Kapton ribbon cables as substrates for the readout chips; this assembly (including an intermediate Be shield and insulator sandwich) is mounted directly on the detector with a minimum of additional mass. A conceptual module layout is shown in Fig. 4-24. A cross section of the electronics-detector sandwich is shown in Fig. 4-25. Note that the Be shield also serves the important function of providing a good thermal path from the readout chips to the cooling ring. This reduces the heat

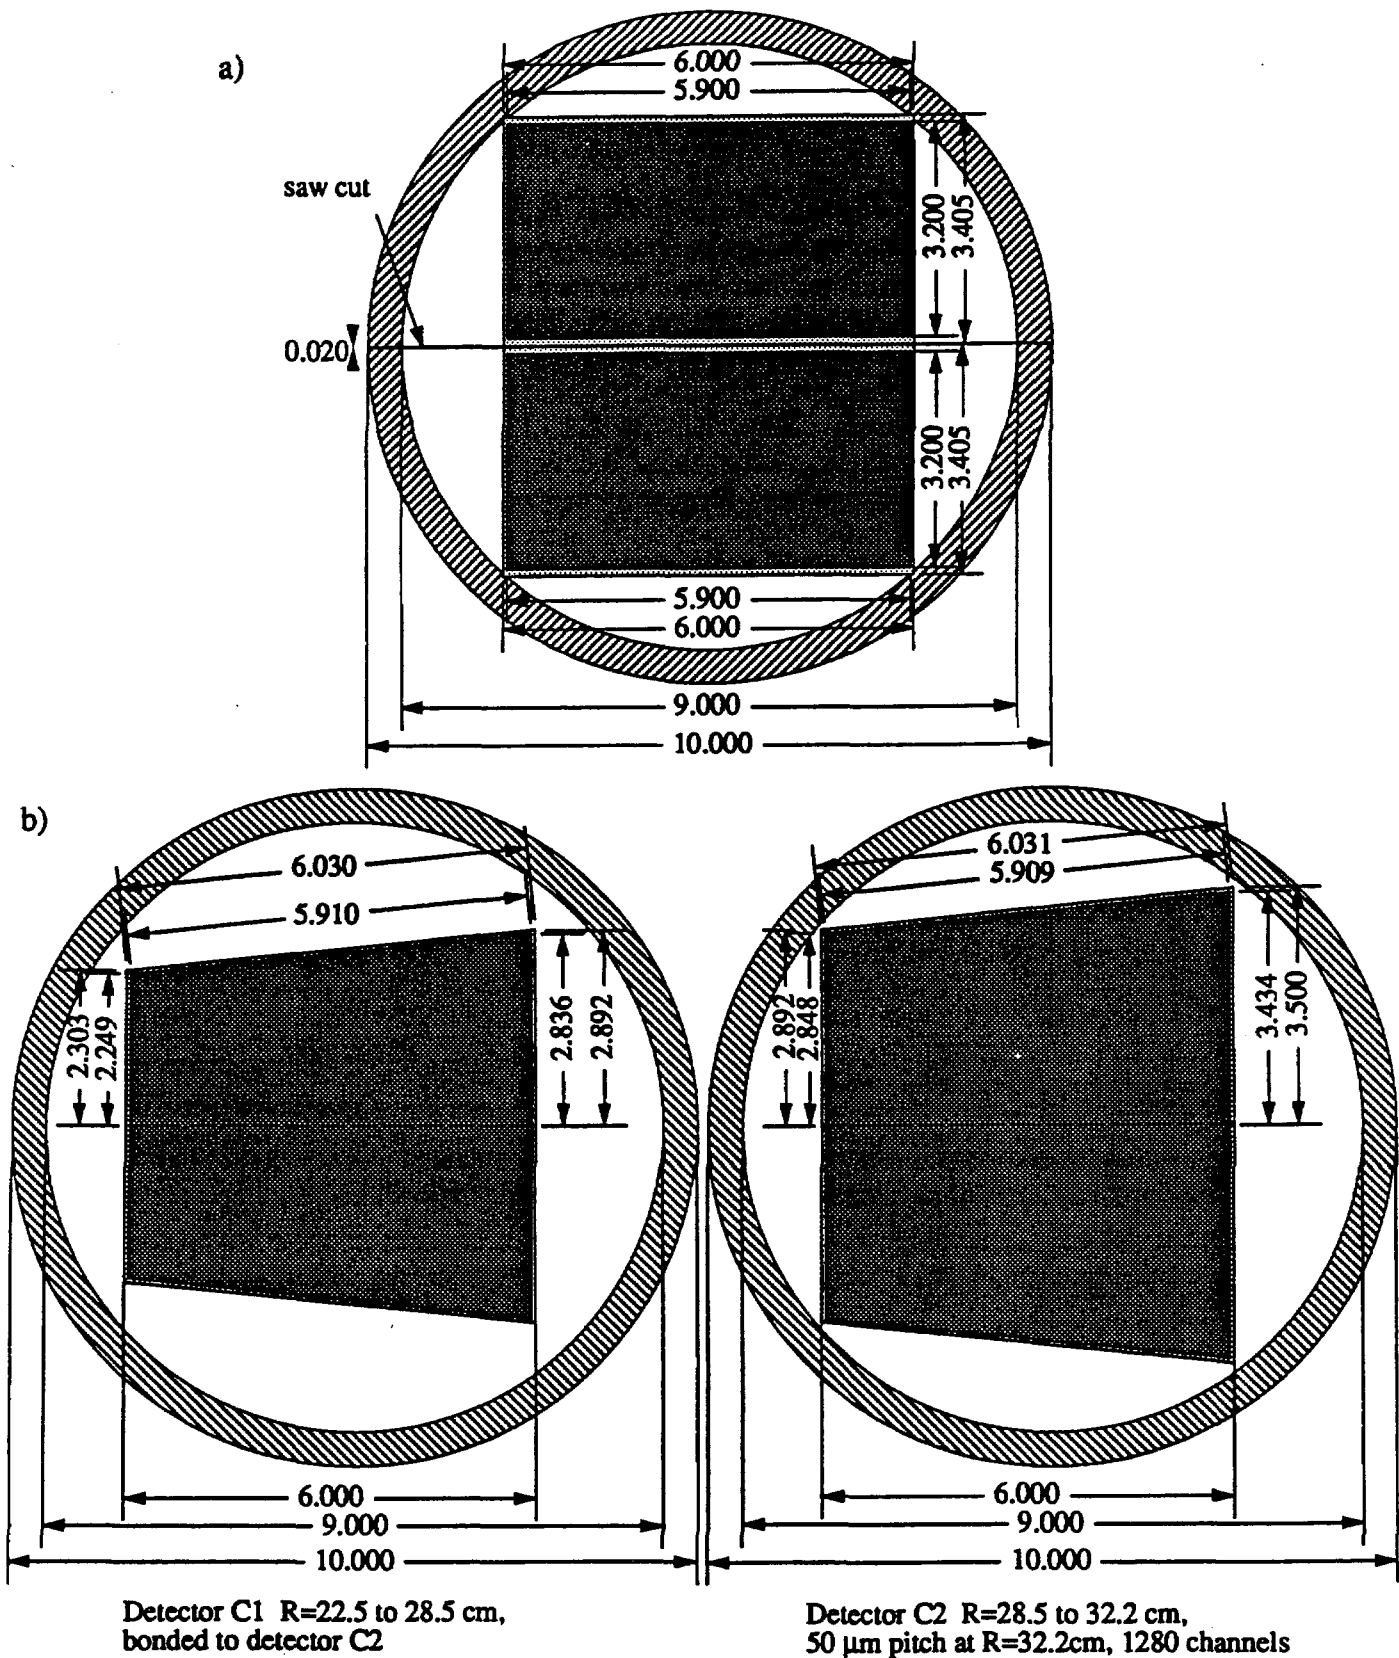


FIG. 4-21. (a) Configuration of two barrel detectors made from one wafer; (b) Geometry of largest disk detector as configured on a wafer. Also shown is geometry of a detector which will be bonded to it.

Table 4-9
Detector and channel count for the silicon tracking system.

Barrel				
Barrel #	Modules in ϕ	Active overlap μm	6cm long Wafers in layer	Electronic channels
1	18	545.6	180	115200
2	24	429.4	240	153600
3	36	347.2	360	230400
4	42	329.8	420	268800
5	48	318.6	480	307200
6	54	310.9	540	345600
7	66	301.3	660	422400
8	72	298.2	720	460800
		Sum	3600	2304000

Disk				
Disk #	Module A in 1 disk (Active overlap 806 μm)	Module B in 1 disk (Active overlap 837 μm)	Module C in 1 disk (Active overlap 775 μm)	Module D in 1 disk (Active overlap 740 μm)
1	26	38		
2	26	38		
3	26	38		
4	26	38		
5	26	38		
6	26	38		
7	26	38		
8	26	38		
9	26	38		
10			32	44
11			32*	44
12				44
13				44*
Electronic channels		1198080	327680	901120
Channels total, all disks				4177920

* Module made of 6 cm units, instead of 12 cm. All others 12 cm.

Table 4-10
Specifications of double-sided silicon strip detector (barrel part).

1) Substrate	
Type:	n-type
Resistivity:	4–8 k Ω ·cm
Thickness:	300 \pm 10 μ m
2) Size	
Overall dimension:	60 mm \times 34.1 mm
Effective area:	58.8 mm \times 32.0 mm
Dead area:	600 μ m from edge
3) Strip	
Pitch:	50 μ m on both surfaces
Strip isolation of ohmic side:	p ⁺ blocking line method
Pattern accuracy	
Position:	\pm 1 μ m
Size:	\pm 1 μ m
Relative position of both sides:	\leq 5 μ m
4) Bias resistor	
Poly-crystalline silicon line on both surfaces	
Resistance value:	250 \pm 50 k Ω
5) Electric properties	
Initial leakage current:	\leq 1 μ A (overall) \leq 100 nA/channel
Bias voltage tolerance:	\leq 150 V
Decoupling capacitance of strip	
Breakdown voltage:	\leq 100 V
Capacitance:	\geq 30 pF/cm
Readout capacitance:	\leq 1.4 pF (junction side) \leq 1.8 pF (ohmic side)
6) Fiducial mark for integration	
Position accuracy relative to strips:	\pm 1 μ m
7) Dicing	
Full cutting by diamond saw	
Cutting zone:	\pm 30 μ m

Table 4-11
Dimensions for various disk modules.

Detector type	Active width at R Out (cm)	Pitch In (μm)	Pitch Out (μm)
A	6.64	28.68	51.63
B	6.55	35.34	51.05
C	6.87	34.94	53.57
D	6.73	38.94	52.48




FIG. 4-22. Structure of each detector surface.




FIG. 4-23. Detector module construction.

path through the detector to the area subtending the cooling ring. Since about 85% of the detector extends beyond the cooling ring away from the readout, heating of the detector is minimized. Thermal analysis shows that the detector temperature is increased by only about 3° C above the cooling ring temperature of 0° C on the end near the heat source. The maximum mechanical deflection due to heating is $< 2 \mu\text{m}$.

Since the Silicon Tracker has several million channels of electronics, efficient on-chip data sparsification is crucial. With only a hit/no-hit indication, one is essentially "flying blind" with respect to signal diagnostics. Hence, reliable control of cross-talk and any other causes of spurious hits is of extreme importance. Cross-talk from cables can be reduced by proper selection of the front-end shaping time constants to reduce electronics response to the spectral distribution of bus activity (especially high-frequency clocks). Differential lines are required to reduce coupling from the cable to the input amplifiers through the detector strips. For local bussing the signal levels on these lines can be kept small; just how small must be determined by careful analysis and experiments.

Decoupling capacitors are potentially an additional significant contribution to mass; circuit and bus designs are under investigation to minimize the number and size of capacitors, or possibly eliminate them altogether. This is being addressed at two levels:

a) Circuit design

The analog circuitry is specifically designed for noise immunity and small transient currents. The comparator threshold voltage is brought in differentially. Signal transmission in the front-end chain from the analog IC to the digital IC is by a localized loop that allows separation of the analog and digital grounds.

b) Power bus design

All power supply lines are configured as low-impedance transmission lines to present a low-impedance return path over a wide frequency range. The power bussing is subdivided to isolate sensitive circuit blocks. The detector bias supply lines are isolated from the front-end chip supply lines (the on-chip referencing preserves common-mode rejection). Analog and digital supply lines (and grounds) are segregated.

The power lines and associated front-end controls are organized by layer in the barrel and by ring in the disk region. Each layer or ring is subdivided in ϕ into 8–12 segments, each of which has a dedicated power cable that is not electrically connected to any other part of the system until the cable terminates outside the detector at a distribution unit. This isolation reduces cross-coupling and limits the fraction of the detector that would be affected by a segment or cable defect. Segmentation also serves to distribute the mass of the power and signal busses more uniformly around the detector. This scheme results in 160 sets of power lines emanating from each end of the detector. Thus a problem in a section serviced by one cable would result in losing 0.3% of the system.

4.4.2. Pixel detector option

High quality track detection close to the beampipe is very desirable for definition of primary and secondary vertices and improved pattern recognition. At the same time, the very high radiation levels and track densities close to the beams create a difficult technical challenge. Pixel detectors based on contemporary hybrid silicon technology could provide the needed tracking capability with radiation hardness superior to strip detectors. We consider this as an option for the inner two layers of the STS, either at turnon or as an upgrade depending on the evolution of these devices, availability of extra funding, and further knowledge gained about the physics priorities for the SSC.

The pixel device is envisaged to be configured in a two-layer barrel at radii of 6 cm and 9 cm, with possible end disks to increase the rapidity coverage beyond $|\eta|$ of 2. This should allow for a reduction of the multiple scattering term in the impact parameter error to about 2/3 the value in our baseline design. The resolution in a pixel layer, using analog information, is expected to be about $5 \mu\text{m}$ in the azimuthal coordinate and $50 \mu\text{m}$ to $75 \mu\text{m}$ in the z coordinate. These should translate, given the rest of the SDC baseline tracking system, into errors of about $5 \mu\text{m}$ by $50 \mu\text{m}$ at the beam vertex for very high momentum tracks.




FIG. 4-24. Conceptual layout of a detector module.




FIG. 4-25. Cross section through a detector module.

In a hybrid technology, pixel detectors are made by joining two independently fabricated silicon devices by a matrix of metallic bonds. "Bumps" of gold, indium, and solder are among the possible materials that have been successfully employed. In the present case, one silicon device is a p-i-n detector with the appropriate pixel pattern on the n side, and the other is a radiation hard CMOS VLSI integrated circuit. Separate processing permits optimization of each device. Costs may be greater than with a monolithic approach, but all the required technological elements appear to be mature.

In a monolithic technology, the detector and readout circuitry are fabricated simultaneously within a single wafer, although the wafer may itself be "composite". The advantages of a monolithic technology derive mainly from working with a "chip" that may be more robust physically, less costly ultimately, and which should offer a slightly smaller overall material thickness. In this case, the technology is still developing, and radiation hardness remains to be incorporated and/or demonstrated.

Because the hybrid technology is the only one for which a clear technological path is available, this approach has been emphasized in current engineering effort and schedule development. Nevertheless, new avenues for monolithic integration will continue to be explored. The thickness of the p-i-n detector is envisaged to be between 250 and 300 μm thick. Although the pixel detector can function well with detectors thinner than 250 μm , such thin wafers cause processing difficulties. The readout IC may be thinned after fabrication to 50 μm , to minimize material.

Because of the extended length along z of the luminous region, tracks in the central region can traverse the pixel detector over an angular range of more than $\pm 45^\circ$ in the r - z plane. For the more inclined tracks, the deposited charge will be shared among several pixels along z if the pixel length is short compared to the ≈ 250 μm detector thickness. Sharing charge over many pixels substantially reduces the S/N ratio. From the limited perspective of preserving a high S/N ratio, therefore, there is a "natural" pixel length of about 250 μm along z .

For vertex reconstruction purposes, however, a much smaller dimension is desired. The ideal resolution is on the order of 5 μm rms in both r - ϕ and r - z . Square pixels, $\sim 40 \times 40$ μm^2 , would be nearly ideal. It does not appear feasible with current technology to pack the needed circuitry into such a small area. A choice can be made between large, square pixels, or rectangular pixels of equal area, which can at least provide good resolution in one dimension. For rectangular pixels with short dimension of about 50 μm , the long dimension appears to be approximately the "natural" length mentioned above. Whether a rectangular pixel is the right choice for physics has not yet been demonstrated through simulation. In the following discussion, the pixel cells are taken to be 50 $\mu\text{m} \times 250$ μm , with the long dimension parallel to the beam axis. With this geometry, there are about 2×10^7 pixels in the system, and ~ 0.25 m^2 of active silicon.

At the standard SSC design luminosity, 10^{33} $\text{cm}^{-2}\text{s}^{-1}$, and at a radius of 6 cm from the beams, pixel cells of this area will be struck at about 340 Hz. The pixel array itself is taken to be 256 rows \times 64 columns, giving an active area of 1.28 cm \times 1.6 cm = 2.048 cm^2 , and a total pixel count of 16384. The array as a whole will be struck by primary tracks at a rate of 5.6 MHz. The chance of the array being hit on any given beam crossing is 9%, a rather large probability. This fact has a major impact on circuit architecture, ruling out centrally managed activity, except during valid data readout.

The question of the optimal organization of the readout for pixel arrays for the SSC has led to examination of several different schemes discussed in Refs. 27-31. The development of a full readout chip is the highest priority of the pixel group.

For the pixel system, both detectors and readout IC's must survive radiation doses in excess of 1 Mrad/year. There exist now quite encouraging results from the UPMC process. Needless to say, the pixel detector faces the most severe radiation levels within the tracking system, and much work remains to be done to establish viability beyond reasonable doubt. The effects of detector leakage current are expected to be very small for the pixel detector, even at very large fluences. The pixel array, by its large segmentation, partitions the leakage current to tiny values. Because of this, the detector may not even need to be AC coupled.

4.4.3. OTD design and modularity

The barrel outer tracking system consists of three axial (trigger) and two stereo superlayers coaxial with the beam direction. A portion of an axial superlayer is shown in Fig. 4-26. The superlayers are made up of close packed trapezoidal modules, with each trigger module containing 212 drift cells, and each stereo module containing 159 drift cells. The maximum length of a module is about 4 m (see Table 4-3). The cell cathodes are formed from light-weight plastic straws, 4 mm in diameter.




FIG. 4-26. Axial straw superlayer modules.

The module concept evolved as a method for assembling over one hundred thousand drift chamber cells into a tracking system. This requires mass production techniques, quality control, and repairability, as well as low cost. In the module concept the straws are grouped into six or eight-layer packages with strong outer covers and their own electrical and gas connections. This manner of construction allows the modules to be mass-produced and tested before they become part of the superlayer system.

Small drift cell structures have found many applications especially in small high precision chambers. There are some important advantages of the tube structure over a cell design with wire cathodes. One is its greater ease of construction and the greater isolation between cells both electrically and mechanically. Also the amount of wire tension which needs to be supported is reduced by a large factor because there are no cathode wires. These considerations were discussed thoroughly in the 1986 Snowmass report [32].

In the 4 mm straw tube, the electric field strength at the cathode is already as high as 2.3 kV/cm, at an applied voltage of 2 kV, and increases to about 45 kV/cm at a point 100 μm from the wire where the avalanche starts. A fast gas is chosen to minimize drift time. The best gas appears to be CF_4 with an admixture of Isobutane (20%) whose drift velocity is $105 \pm 15 \mu\text{m/ns}$ at fields up to about 30 kV/cm and somewhat larger at higher fields. A straw tube with a diameter of 4 mm operated with CF_4 is compatible with the 16 ns bunch crossing time in the sense that the electrons from the passage of charged particles through the tube in one bunch crossing will mostly be collected before the next bunch crossing occurs. The maximum drift times in a 2 T field are expected to be about 29 ns. Thus all charge will be collected in two crossings.

The gain must be as low as possible consistent with obtaining good resolution so that the current draw is minimized. A gain of 2×10^4 will keep aging effects small. The resolution is related to the gain and the gas properties. The resolution has been measured to be about 100 μm in a short section of a straw tube assembly where only the performance of the gas and electronics is tested and not the mechanical stability and precision of the structure. These studies are being extended to larger modules and better electronics. Prototype chambers have been run at up to 1 MHz with no apparent effect on the resolution. The following table indicates the expected occupancies and current draw per unit length for the five superlayers in the system at the design luminosity. The effects of neutrons have been calculated for each layer and added to the current draw expected from charged particles.

Table 4-12
Occupancy and current per tube at design luminosity.

Superlayer	Radius (m)	Length (m)	Occupancy	I (Charged part.) nA/cm	I (Total) nA/cm
1	0.82	2.41	0.096	1.7	1.9
2	1.10	3.38	0.063	0.72	0.9
3	1.35	4.03	0.043	0.39	0.6
4	1.49	4.03	0.035	0.32	0.5
5	1.63	4.03	0.026	0.24	0.5

Design of long chambers

Straw modules have been constructed with lengths up to 4 meters. To keep the anode wire stable, it will have to be supported every 0.8 m along its length. A wire support design has been developed. Signal attenuation, another important consideration in a long chamber, arises almost entirely from resistive losses in the anode wire and the cathode coating. The attenuation length is given by $2Z_0/R_L$, where Z_0 is the characteristic impedance of the tube and R_L is the sum of the anode and the cathode resistances per unit length. A cathode thickness of 0.15 μm of copper leads to a cathode resistance of 17.5 Ω/m of tube length. With a 38 μm gold plated tungsten anode wire an attenuation length of 6.8 m is expected, which leads to a 44% attenuation for a pulse traveling the full 4 m length of a tube.

A terminator is required at the far end of the tube from the amplifier to terminate the tube transmission line in a resistor equal to the characteristic impedance of the line. A capacitor is also needed for blocking the DC voltage and for compensating for the lossy nature of the transmission line. The characteristic impedance of a transmission line formed by a 38 μm diameter inner conductor and a 4 mm diameter outer conductor is 279 Ω . The optimum capacitor size for a tungsten wire and a 0.15 μm copper cathode is about 90 pF.

Module design

The design of the graphite/epoxy straw module shell is shown in Fig. 4-27. Each side of the foam panel is covered with a 2 ply carbon fiber layer. By using a 0.005 cm fiber ply and a Rohacell IG51 foam the shell thickness is only 0.36% X_0 . Figure 4-28 shows a detail of the end of a module. The straws and wire supports are premounted in the shell. The wire supports are attached inside the straws and near each end. The shell clamps the straws and wire supports in a closed-packed geometry. The endplate acts as a tension plate for the wires and as one side of a gas manifold. The printed circuit board for the electronics is attached to the endplate. This construction results in a very short end section on each module and a very low mass connection. Our goal is to keep the total material to a few percent of a radiation length.

Layout of trigger modules

The layout of the modules has been optimized for the five superlayer tracker. A program to optimize the size, position and separation of each module in a superlayer was developed. In particular the study investigated ways that the modules could be positioned so that high momentum particles have a minimum number of lost hits in the boundary regions of modules. One way that this can be accomplished is to stagger alternate modules in the radial direction. This radial stagger reduces the space between modules and maintains a maximum number of hit straws. Configurations have been found that will allow high momentum tracks to have the full complement of hits. A typical layout is shown in Fig. 4-29.




FIG. 4-27. Straw module shell.




FIG. 4-28. Detail of end of module.




FIG. 4-29. Axial module overlap region.




FIG. 4-30. Module positioning for stereo straw layers.

Layout of stereo modules

The tracker design incorporates six-layer stereo modules in two of the outer superlayers. Each one is positioned at about a 3° angle with respect to the cylinder axis. The separation of the adjacent modules thus changes as a function of the distance from the center of the cylinder. These modules are also positioned at two different average radial distances, so that there is complete coverage for all tracks in the overlap region of adjacent modules. The position of the modules in this stereo configuration has been worked out in detail and is shown in Fig. 4-30.

4.4.4. ITD design and module

As discussed in the introduction, the ITD is segmented into bites of approximately 0.25 units in η and made of annuli composed of gas microstrip tiles. The dimensions are listed in Table 4-4. Including overlap, tiles of length up to 270 mm are required, which falls within the manufacturing limits that we confidently expect. The tile widths are chosen not to exceed the manufacturing limitation, which is an active width of 180 mm, and to be identical within each annulus thus minimizing manufacturing costs. A typical gas microstrip tile structure is shown in Fig. 4-31.




FIG. 4-31. Typical gas microstrip module structure.

The microstrip tiles (GMDs) consist of a microstrip substrate and a plane of drift electrodes separated by a small gap of about 3 mm. Both the substrate and the drift plane will require support to enhance their rigidity. The edges of the tile are sealed to create a gas volume, gas connections being provided to allow gas to be circulated through the drift volume. The anode and cathode lines are laid out radially on the substrate. HV connections to the cathode lines are made at one edge of the tile, the signal connections to

the anodes on the opposite edge. It is also necessary to make a HV connection to the drift electrode and to an electrode on the back of the microstrip substrate.

The anode strips on the GMDs run radially from the beam axis. As a result, the distance between the anode strips increases with radius. If the widths of the anode strips and the size of the gap between anodes and cathodes is kept constant, the width of the cathode strips can increase by more than a factor of two over the radial dimensions of a GMD tile. We call this "keystoned" cathode geometry. Investigation of the consequences of keystone geometry on the variation of gain over the anode length has been carried out for two chamber geometries which differed only in the widths of the cathodes.

The program MSFIELD [33] was used for these calculations. The program calculates the gain by integrating the Townsend coefficient over the drift paths of electrons arriving at the surface of the anode. We find that the effect of a factor of two keystoning results in a predicted 60% increase in gain. Work has been in progress to arrive at optimum average operating conditions for the chambers, and to tune the parameters of anode width, cathode width, anode-cathode spacing and voltage settings to reduce this effect. Recently, we have been able to arrive at a keystone configuration shown in Fig. 4-32, where the anode width also increases with radius. To keep the same anode to cathode gap, the cathode width increase is reduced correspondingly. Gain stability of a few percent over the whole anode length appears to be possible with a reasonable overall optimized voltage configuration. A series of GMDs will be made with varying pitches to study the effect of keystone geometry on GMD gains and verify our optimized design.




FIG. 4-32. Geometry of ITD electrodes.

Studies of drift velocities and gain characteristics of gases suitable for the ITD environment will be carried out using laser test cells and sources. A desirable gas mixture should provide a gain of a few thousand, which maintains a good signal-to-noise ratio, with minimum aging effects, and maximum drift velocity. The Xe/DME/CO₂ and DME/CO₂ mixtures and various CF₄ based gases will be studied. The CF₄ based mixtures are known to have good aging properties. Their gain may be lower than other gases but may be adequate for ITD applications.

The front-end electronics are mounted on a board which is rigidly attached to the GMD tile. The exact location and shape of this board will depend on the method used for bonding the anode microstrips to the front end electronics. The electronics require LV power and cooling. Two or three optical fibers per tile will be used to transmit the data during data acquisition, and up to 40 optical fibers per tile for the radial anodes used in the Level 1 trigger.

The GMD-ITD provides a very low occupancy system. Simulations indicate that the occupancy will be at most 0.4% per channel at design luminosity. This assumes that soft looping tracks trigger all cells crossed independent of path length which provides an overestimate.

Initial work on the tiles themselves has so far been restricted to development of prototypes for physics testing and evaluation. Glass, plastic, and silicon substrates are being investigated. Detailed design work on the mechanical support for the tiles will depend greatly on the result of the R&D effort. In the meantime, generic problems such as connections, connectors, cathode support, tile mounting and tile location are being investigated.

4.5. Electronics

The front-end electronics for the silicon strips and gas microstrips are bonded directly to the detectors and are a highly integrated part of the detector module. For these systems we discuss the front-end electronics in this part of the technical proposal. For the straw tube detectors, the detailed front-end electronics design is discussed in section 8.4 of this report.

4.5.1. Silicon front-end electronics

Each silicon detector channel contains the following functions:

- a) low-noise preamplification, pulse shaping and discrimination provided on a bipolar chip;
- b) time stamping, data buffering and sparse readout provided on a CMOS data storage chip.

The two chips are bonded together to form the full front end. Both bipolar and CMOS chips process signals from 128 strips of 50 μm pitch. The chips are expected to be about 4 mm and 6 mm long, respectively. The mounting arrangement has been discussed in section 4.4.1.

A simplified circuit diagram for the bipolar circuit is shown in Fig. 4-33. Each channel comprises a charge-sensitive preamplifier followed by gain stages and a comparator. CR-RC pulse shaping is implemented by the discharge time constant $RfCf$ in the preamplifier feedback loop and by an RC integrator formed by $R11$, $R12$, and $C11$ at the output of the differential gain stage $Q11$ - $Q12$. The threshold is applied differentially to $Q13$ and $Q14$ to avoid threshold variations introduced by noise or voltage drops on the supply lines. The reference voltage applied to the base of $Q12$ is provided by a "dummy" amplifier identical to the stages feeding $Q11$. One "dummy" amplifier serves an entire chip. The circuit is easily adaptable to either positive or negative input signals (as required for double-sided detectors), by juxtaposing the output connections at $Q11$ - $Q12$; or, since the circuit will handle bipolar signals, by inverting the digital output. The output stage is shown as a current driver, which can be scaled easily to match the drive requirements of the following chip.

The power dissipation per channel for the bipolar chip is 620 μW . The shaping time constant (CR-RC) is 20 ns and the supply voltage is 3.5 V. Some of the specifications for the front end are as follows:

- a) 1200 electrons noise for the expected detector capacitance of 14 pF.
- b) Less than 15 ns time variation for $1\text{ fC} < Q(\text{signal}) < 8\text{ fC}$.
- c) Dead time $< 100\text{ ns}$ after a hit.
- d) Less than 0.5% inefficiency in a channel.

The data from each strip are buffered on the CMOS chip at every beam crossing for a period of time called the Level 1 latency which is 256 beam crossings. During the Level 1 latency, a Level 1 trigger decision is made either to reject the data from the corresponding beam crossing or to retain it for transmission off

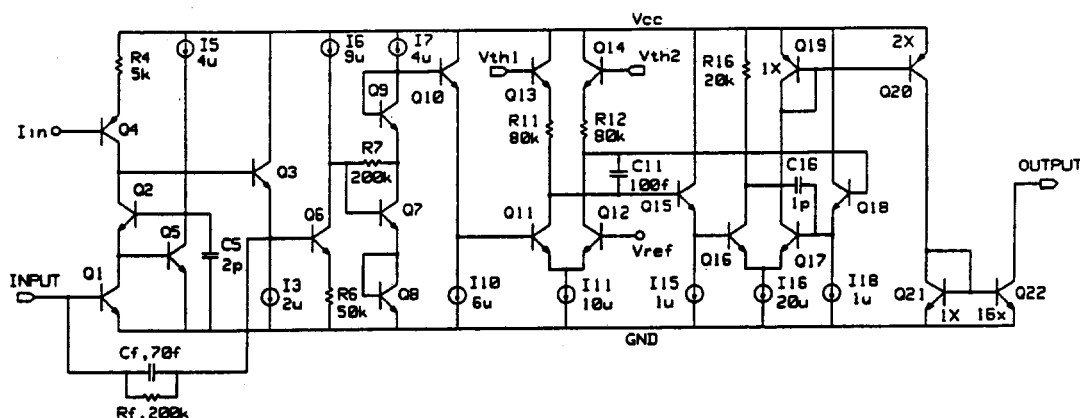


FIG. 4-33. Simplified circuit diagram of the analog front-end IC.

of the chip. The data from a beam crossing will be channel address encoded once a decision is made to retain the data from that beam crossing. The encoded data will then be serially transmitted off of the chip.

The channel address encoding is data-driven so that the chip only stores addresses of the channels that received hits. The encoded data from the beam crossings that received Level 1 triggers will be transmitted in the same order as the Level 1 triggers.

Figure 4-34 shows an abstract representation of the data-driven readout architecture. In this scheme there is a CAM (content addressable memory) ARRAY assigned to each channel. Each array contains five locations, determined by the necessity to keep the efficiency high. This allows an efficiency $>99.5\%$ even for 0.5% occupancy. Each location is capable of storing the time stamp generated by the time stamp counter and also indicating at every beam crossing whether the contents match the desired time stamp. The length of the time stamp sequence is 8 bits for 256 crossings, and it is synchronized to the beam crossing. Therefore, within every time stamp sequence, each time stamp corresponds to a beam crossing.

If there is a hit at a beam crossing in a given channel, i.e., its input is active, then the time stamp of this beam crossing will be stored into a location for this channel. After the Level 1 latency the time stamp will be the same as the stored time stamp. At this point if there is a Level 1 trigger (which will correspond to the beam crossing at which the time stamp was stored) this channel will output a MATCH signal. Therefore this scheme enables the data at the input to a chip to be reconstructed after Level 1 latency by grouping MATCH outputs from all the channels.

The power required for the digital data storage is expected to be about $200 \mu\text{W}$ per channel.

4.5.2. ITD electronics

The design of the ITD front-end electronics is determined by the signal characteristics of the GMD and by the trigger and data taking requirements.

GMDs are usually operated at fairly low gas gain and for the basis of this discussion, we assume a gas gain of 5×10^3 which is typical of existing devices. The track length for a particle traversing the GMD varies between a maximum of 3 mm, when the track is perpendicular to the GMD tile, to a minimum of $200\text{--}400 \mu\text{m}$ per cell, when the track is parallel. High momentum tracks will deposit ionization in the full 3 mm of gas. Because this is a relatively small path length compared to conventional proportional chambers and because we intend to operate these GMDs at atmospheric pressure, a gas with good ionizing properties is needed. We assume the use of a gas similar in ionization properties to Xenon or DME, and producing 200 ion pairs per cm. In that case, the signal after gas gain will be $3 \times 10^5 e$. For a standard drift chamber gas with a drift velocity of $50 \mu\text{m}/\text{ns}$, electron collection times of up to 60 ns are expected. Since risetimes of 20 ns or better are needed, the useful collected charge is $1 \times 10^5 e$.

These numbers are subject to several other effects that can degrade the pulse height:




FIG. 4-34. Outline of the data-driven readout.

- Landau fluctuations can lead to a minimum energy deposit which can be as low as half of the expected value;
- tracks can divide the charge between adjacent anodes;
- fluctuations from primary ion statistics can yield large pulse height differences.

These phenomena have been looked at and an evaluation made of the expected timing fluctuations for a GMD plus front-end. We assume for this evaluation the following circuit specifications:

- 1000 e noise, based on 6 pF detector capacitance,
- 20 ns risetime,
- 3.5 σ of the noise for a discriminator threshold.

These lead to fluctuations in trigger timing of

- < 7 ns due to ionization statistics,
- < 5 ns due to noise and time slewing.

These can be taken in quadrature and yield a time jitter of < 10 ns. Thus, these devices are capable producing a Level 1 trigger.

The GMD front end is expected to be a two chip set rather similar to that of the silicon. Differences are primarily the need to provide a separate information stream for the Level 1 trigger and the different pitch of strips on the GMDs. We expect 32 strips to be read out by one chip, and that the digital data storage, which is read out only after a valid Level 1 trigger and contributes information at Level 2, can be organized identically to the silicon data storage. The readout for the ITD is therefore organized in a Level 1 trigger readout path and the data storage path. Both use fiber optics. The optimum level compaction and multiplexing for both readout chains is presently under study.

4.5.3. Straw tube electronics

Design requirements

For each straw drift tube it is necessary to accurately measure the time of arrival of the first electron (or cluster of electrons) at the anode. This, combined with the desire to operate with as low a gas gain as possible, implies the use of a low noise preamplifier with a risetime sufficiently fast to provide the desired time resolution, but sufficiently slow to provide acceptable signal-to-noise. In addition, because of the high rate of pulses on individual wires excellent double pulse resolution is very important.

The design goals are:

- a) < 0.5 ns time accuracy in order to ensure that the contribution of the electronics to the spatial resolution is small compared with the 100 μm intrinsic resolution.
- b) Double pulse resolution of 20–30 ns. We have adopted the specific goal of having the return baseline for a single cluster occur in less than 15 ns.
- c) Semi-Gaussian shaping to minimize baseline shifts and noise from parallel current sources.
- d) On-chip Level 1 storage for 4 μs and Level 2 storage for a latency interval of order 100 μs .
- e) Ability to simultaneously acquire new data, store data, and read out interesting data without self-interference or interference to other systems.
- f) Ability to withstand $\sim 1 \text{ Mrad}$ and $> 10^{14} \text{ n/cm}^2$ over the life of the experiment.

A more complete description of the actual electronic implementation is contained in section 8 and this section is mainly limited to a description of the mechanical aspects of the detector mount electronics—space required, mounting, connections, and cooling for the front-end board (FEB). Note that the term FEB encompasses all of the elements located at a module endplate on the detector—HV system, active substrate, and time substrates.

Connection to the detector

Both because of the tight spatial constraints in the tracking volume and because of the possibility of electrical pickup and interference, it is important that the connections to each of the anodes and cathodes be as compact as possible. Nevertheless, in order to allow assembly of the detector and electronics, and to facilitate testing and maintenance, it is important to have an electrical and mechanical disconnect at each endplate of the straw tube modules. The details of this disconnect are not yet complete, but it could be based upon either standard pin and socket connections or sheet elastomeric connections. A pin/socket connector of more than 200 contacts has the advantage that it would be nearly self supporting and thus would not require large additional mechanical fasteners to hold the electronics on to the endplate. Such a connector, however, would imply significant additional mass and thus the low density elastomeric solution is also being investigated.

The part of the FEB directly in contact with the endplate is the High Voltage distribution and coupling network. The HV coupling capacitors are mounted between the first and second HV Substrates (*substrate* rather than *printed circuit board* because a wide variety of plastics, ceramics, and even glasses are under investigation) as sketched in Fig. 4-35. It is not yet clear whether or not the stereo layers will need different HV Substrate designs because of the nonuniform placement of the modules.




FIG. 4-35. High Voltage Substrate design—elastomeric connector version with a side view and end view of a module.

Active electronics

The front-end electronics are mounted on a 10 cm square board that is attached to the HV substrate. Four rows of bipolar front-end chips (ASD) and two rows of trigger chips are mounted directly on the front and back of the active substrate in high density (lead spacing about 0.5 mm) packages. The data collection chips (DCC), the cable to the crate interface card, and assorted passive objects are mounted on the back of the active substrate. The time measurement and event buffering are located on time substrates mounted perpendicularly off the front of the active substrate. Each time substrate carries the time measuring chips for 16 channels plus decoupling capacitors.

The ASD, trigger, and DCC chips are separately packaged in high density packages that allow removal and replacement of any individual defective chip. The advantage of repair seems, at present, to outweigh the disadvantages inherent in very dense, nonstandard packages. The time function is split into several small time-substrate boards to allow repair and replacement of a small number of devices if one should fail. The chips on these boards would be mounted directly to the time-substrates with only an overall protective coating, and without individual packages.

The total power dissipation per channel on the detector will be about 30 mW. This is divided up about evenly between the active and time substrates, and thus cooling an individual front-end board cannot be a simple matter of cold plates—the mechanical shape is too complex. The most suitable method of cooling seems to be airflow across the FEB structures (front and back) with heat exchangers located outside of the tracking volume.

At 30 mW per channel, a trigger FEB will dissipate about 6.4 W and a stereo FEB will dissipate about 4.8 W. Individual superlayers will have from 200 to 600 W per end to dissipate as can be seen in Table 4-13. The total power per end is less than 2 kW and for an air temperature rise of 20°C, there would need to be about 80 L/s. Even if we were to demand a very restrictive 10°C maximum rise, the total volume is still only a few muffin fans worth per end. In a realistic final mechanical design, the air flow paths will not be trivial, and small fans will not be sufficient, but the cooling problem should be tractable.

In the present design we are assuming that power regulation takes place at the crates, not on the FEBs. As a consequence of this assumption, significant power is *not* dissipated in the tracking volume, but we will

Table 4-13
Front-end boards and power load per end of the detector. The air flow in liters/s for a 20°C temperature rise is calculated from a Rotron rule— $Q(\text{L/s}) = 0.83 \times \text{Power}/\Delta T(^{\circ}\text{C})$.

Superlayer	Trigger/stereo	FEBs	Power/end	L/s
1	Trigger	46	292	12.1
2	Stereo	62	296	12.3
3	Trigger	76	484	20.8
4	Stereo	84	400	16.6
5	Trigger	92	585	24.9

have to accept a significant voltage drop along the power feed lines into the tracker. If this voltage drop becomes too difficult, it may be necessary to put regulation on the FEBs—this would raise the total power dissipation (and thus the cooling load) by about 30 to 50%.

4.6. Mechanical systems and utilities

4.6.1. STS mechanical design

The silicon tracking system (STS) is composed of silicon strip detectors arranged both in a cylindrical array (central region) and an array of flat panels about the interaction point. Fig. 4-36 depicts the silicon arrays without structural support. The overall length of the silicon array is 5.16 m and the maximum diameter is 0.93 m. The Silicon Tracking System Conceptual Design Report [3] should be consulted for the body of analysis performed to quantify the present design concept. For the STS to achieve its physics goals, the mechanical structures and services must:

- support 17 m² of silicon detectors and stabilize their positions at the 5 μm level,
- uniformly cool the detector system to 0° C and at the same time remove 6.5 kW of waste heat generated by the detector electronics,
- provide 1700 A of current to supply the 6.5 million electronic channels, and
- supply all the control and data transmission lines for those channels.

These nontrivial objectives must be attained in a high ionizing radiation environment, using minimal mass and only low-*Z* materials, and comprise a maintainable system for the more than 10 year operating life of the system.

Figure 4-37 depicts a cross-section of the silicon tracker and identifies the major mechanical components. The method for supporting the individual central region silicon modules combines adjacent layers into shell assemblies. Three silicon shell assemblies, two 24-cm and one 12-cm long, comprise a 60-cm-long central region superlayer. The four superlayers in the central region are supported pairwise by two composite support cylinders made of graphite/cyanate-ester resin (G/C-E). This support arrangement minimizes structural material. Figure 4-38 depicts one of the 24-cm-long silicon shell assemblies and demonstrates the hermeticity around the circumference of each silicon layer that is provided by the edge to edge overlap. Module specifications are listed in Table 4-14.

The silicon modules are bonded to structural rings made from G/C-E to form the silicon shell assembly. These rings, hereafter denoted cooling rings, in addition to providing structural support, provide a passage for an internal heat pipe wick artery structure. The detectors in the central region are canted at an angle of 7.4°, therefore the top and bottom surfaces of the cooling ring take on a serrated appearance. A hydrocarbon working fluid is wicked to this area. Heat from the electronics flows through the thin-walled




FIG. 4-36. STS detector arrays (pictorial view).

composite cooling ring surface, onto which the module is mounted, and evaporates the cooling fluid medium. The baseline design uses butane as the working fluid.

The two forward regions are each composed of thirteen silicon annular disks. The detector modules in the forward region are not canted, which significantly simplifies the cooling ring design. Large separations between individual silicon layers make the prospect of combining adjacent layers into superlayers unattractive. Indeed, only for silicon layers near the central region are layers in sufficient proximity for this arrangement to be feasible. Individual modules are arranged circumferentially around the beam axis and are approximately twice as wide as those in the central region. The forward modules are in the shape of a trapezoid in order to fit as individual arc sectors. Circumferential hermeticity is preserved by overlapping adjacent trapezoidal modules. Further simplification of the cooling and structural ring construction is provided by having adjacent modules in one disk on alternating sides of the cooling ring as depicted in Fig 4-39. A more complete description of this geometry is discussed in a subsequent section.

The technique for cooling the forward region silicon modules will be evaporative as previously discussed. Cooling and structural rings which support the trapezoidal-shaped modules must track the silicon wafer contraction or they may introduce distortion. The G/C-E material properties are tailored to match that of silicon, and the module support interface will have built-in compliance to allow for a minor material coefficient of thermal expansion (CTE) mismatch.

A space frame network has been configured to provide support for both the central and forward regions. A metal matrix composite (MMC) material, tailored to exhibit a zero CTE, was chosen for the purpose of eliminating distortion due to thermal and/or radiation effects. The tubular frame will be assembled with demountable joints which are adhesively bonded. Adhesives are used to eliminate the possibility of creep in mechanically fastened joints which are subject to temperature cycling.

The entire silicon tracking system is enclosed in a vessel that is configured to safely contain the butane vapor of the cooling system. Introduction and removal of butane into and out of the containment vessel




FIG. 4-37. STS layout (side view).




FIG. 4-38. Silicon detector module assembly demonstrating particle hermeticity. Shaded areas show, from top to bottom, a single wafer, a detector unit, and a detector module.




FIG. 4-39. Forward region planar arrays demonstrating trapezoidal-shaped detectors. Also shown are radial webs connecting the two cooling rings.

Table 4-14
Silicon detector module specifications.

	Central region	Forward region
Shape:	Rectangular	Trapezoidal
Active detector width:	3.2 cm	~6.4 cm
Number of channels:	1280 (640 per side)	2560 (1280 per side)
Heat dissipation:	1 mW per channel	1 mW per channel
Module heat load:	1.28 W	2.56 W
Thermal conductive area: (at cooling ring surface)	0.8 cm × 3.3 cm	~0.8 cm × 6.5 cm
Module heat flux: (at cooling ring surface)	0.5 W/cm ²	0.5 W/cm ²
Operational temperature: (at silicon detector surface)	0° C	0° C
Cant angle:	7.4°	0°

will be accomplished in a manner imposing only slight differential pressures (0.1–0.2 atms) across the containment vessel wall. Any pressure differential excursion outside the 0.2 atm limit will be automatically vented under controlled conditions.

Kinematic mounts will be used to support the internal space frame to the outer containment vessel wall. Since the butane vapor pressure is controlled by the condenser pressure, internal pressure excursions are not expected. Slight outside pressure excursions (up to 0.1 atm) will result from barometric pressure variations. At the points where the space frame mounts interface with the containment wall, deflections from these small pressure excursions will be minimized. Nonetheless, the internal kinematic mounting of the space frame will be configured to ensure that the STS axis does not become misaligned to the outer tracker reference frame.

Design concepts chosen to achieve the STS performance goals are presented in the following sections. The design effort has concentrated on addressing the stringent stability, positional accuracy, radiation exposure, and radiation length limits required for this high resolution tracking system.

Materials

The material selection for the silicon tracking system is driven by three primary factors:

- a) 5 μm stability in $r\phi$;
- b) a 3% radiation length restriction; and
- c) radiation resistance in the presence of hydrocarbon fluids.

Nominal detector dimensions of 1 m in diameter and 5 m in z dictate that linear expansion effects be minimized, whether from thermal or moisture effects. Graphite/epoxy composite materials, while possessing tailorable CTEs, are typically hygroscopic, i.e., they are characterized by a coefficient of moisture expansion (CME) that ranges from 300 to 600 ppm for a fully saturated substrate. For large spatial structures this magnitude of expansion would be unacceptable (~3000 μm in 5 m). During the conceptual design effort we have studied material choices at length and have narrowed the options to MMCs, whose CTEs can be tailored to be approximately zero, and hydrophobic cyanate ester resin systems for the G/C-E composite candidate. Although the cyanate ester resin systems have low moisture saturation levels, typically 1% (producing an overall laminate CME that is 1/50 that of a graphite/epoxy laminate), there are still uncertainties in overall stability in a radiation/butane environment. The CTEs of G/C-E composites are also tailorable and these materials have a large radiation length. In short, our design strategy is to judiciously apply these material options in the most cost effective manner while insuring that we achieve the

above design requirements. Attainment of our goals will require careful application and thorough testing of all materials under conditions simulating the actual environment. Furthermore, compatibility of the chosen materials must be demonstrated, *e.g.*, adhesives must not damage the silicon detectors or degrade in the radiation environment. Much of this work remains to be done in greater detail and forms the basis of the preliminary design effort.

Radiation length materials considerations

Throughout our design, the need for ultralightweight, low- Z , radiation resistant materials has been stressed. Table 4-15 lists the components which individually contribute to the overall radiation length. The amount of material vs. rapidity in the baseline design has been shown earlier in Fig. 4-14. The design meets our goals but not comfortably. This is an area of concern, and our preliminary design effort must address how further reductions in support materials may be attained. Clearly, the radiation length requirement has a significant effect on the design approach and structural materials chosen.

Table 4-15
Detector and material parameters.

Component	Material	Radiation length (cm)	Thickness (cm)	Comments
Beam pipe	Beryllium	35.3	0.060	At $r = 3.0$ cm
Inner enclosure vessel	Beryllium	35.3	0.050	At $r = 5.0$ cm
Outer enclosure vessel	G/C-E	24.9	0.071	At $r = 50$ cm, length = 530 cm
Silicon detector	Silicon	9.36	0.0300	Cylinders and disks
Barrel support	G/C-E	24.9	0.071	Between barrel layers 2 and 3 and 6 and 7
Space frame	Mg-MMC	16.30	0.11	Tubes of 2.54 OD in a hexagonal arrangement
Cooling ring	G/C-E	24.9	0.045	Rings of rectangular cross-section, 3 cm long by 0.8 cm wide
Wick	Liquid butane at 50% density	148	0.040	Against three insides of of each cooling ring
Silicon chips	Silicon	9.36	0.0300	Width = 1.0 cm
Kapton substrate	Kapton	28.4	0.0200	Width = 2.6 cm
Beryllium shield	Beryllium	35.3	0.0300	Width = 2.6 cm
Cabling	Beryllium	35.3	Variable	Both power and signal: 0.063 cm at the outer wall of the enclosure vessel, extending from there to the detector layers.
Filler gas	Gaseous butane	16000	Filler	Fills remaining space inside enclosure vessel

Enclosure/supports

The enclosure/support system for the silicon tracker performs two important functions. First, the shell assembly provides containment of the evaporated butane coolant (vapor space). Secondly, the outer shell structure serves as the primary conduit for transferring the mechanical loads on the silicon tracker space frame to the outer tracker space frame. Section 4.6.8 discusses particular design/construction aspects of the STS support method and interfaces with the outer tracker. The outer structure, and the kinematic mounts prevent extraneous loads from being directly imposed on the delicate silicon structures, while maintaining precision alignment. Hence, the outer shell structure is a *key* member in maintaining the stability of the STS.

An ultralightweight G/C-E sandwich wall construction has been chosen for the outer shell and two end cover plates. A very thin beryllium (0.5 mm) shell will be used for the inner liner. This combination will provide a stiff low-mass structure. A description of the individual wall structures is contained in Table 4-16.

Table 4-16
Enclosure description.

Component	Material	Size	Remarks
Outer shell	G/C-E (P75/954-3)	5.6 m long \times 1.2 m diam.	Truss core sandwich wall 3-300- μ m-thick G/C-E lamina
Inner shell	Beryllium	5.6 m long \times 10 cm diam.	0.5 mm wall with reinforcing ribs
End covers	G/C-E (P75/954-3)	1.2 m diam. \times 2.54 cm wall	Double truss core composite 5-300- μ m-thick G/C-E lamina

Under normal operation, the shell structures will not be exposed to a pressure differential, since the butane cooling system works at one atmosphere. Also, filling and emptying the butane system will be accomplished without imposing any significant pressure differential. Consequently, the gravitational loads associated with the various structural mass components largely dictate the stiffness requirements for the outer shell. In contrast, the end cover structural design is dictated by the 0.2 atm maximum design pressure differential chosen.

Access to the internal cable connections at each detector end is provided by the removable end covers. Provisions for access *in situ* will be required. Consequently, the outermost shell/flange combinations must be sufficiently stiff to carry the entire support reactions with the covers removed. With this prospect in mind, we have chosen a sandwich shell wall construction with a truss core possessing a 45° included angle that provides considerable longitudinal stiffness and minimizes gravity sag along the shell axis. The prototype, depicted in Fig. 4-40, consists of a truss core sandwich composite with a 70° included angle. In addition, in the vicinity of the shell flange, the truss core will be circumferentially reinforced to minimize radial deflection when the covers are removed. The flange, a G/C-E honeycomb structure, provides most of the circumferential rigidity necessary for the silicon tracking system mounting interface. Figure 4-41 depicts the location of these reinforcements.




FIG. 4-40. Sandwich construction shell demonstrating sparse, truss core geometry.




FIG. 4-41. Enclosure shell demonstrating end flanges and intermediate reinforcements.

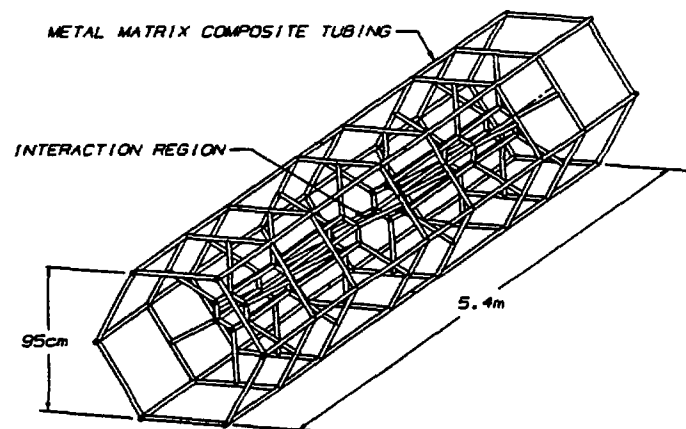


FIG. 4-42. STS space frame composed of MMC tubing and possessing 12 fold symmetry.

Space frame

The space frame provides structural support for the cylindrical arrays of the central region and the flat, planar arrays of the forward regions. Table 4-17 enumerates the space frame design parameters while Fig. 4-42 demonstrates its construction. The dimensional stability requirements motivate the use of MMC construction materials. Their ultrahigh elastic modulus is supplemented by their immunity to moisture-induced distortions and the capability to tailor their longitudinal CTE to approach zero. This materials selection also facilitates the use of thermoplastic adhesives in the demountable joints as a design option. The frame features an inner hexagonal geometry rotated about its longitudinal axis to avoid aligning with the outer hexagonally arranged members. This diminishes the radiation length seen by any one particle. Finally, the MMC space frame readily satisfies the requirement for butane compatibility in a high radiation environment.

Table 4-17
Space frame design parameters.

Length	5.4 m
Effective diameter	1.1 m
Configuration	Hexagon with 12-fold symmetry
Materials	P120s graphite (55 v/o) / Az91 magnesium (45 v/o)
One radiation length	16 cm
Structural members	Circular cross-section tubes
Member dimensions	25 mm OD (1.1 mm thick)
Frame mass	~25 kg
Longitudinal modulus	~470 GPa
Longitudinal CTE	~-0.5 ppm/°C
Mounting scheme	4 points (kinematic)
Fundamental frequency	~40 Hz

Preliminary structural analyses have been performed to aid the design process. Several significant issues must still be addressed during the preliminary design phase. These include structural testing of the MMC tubes and correlation of these results with finite element (FE) predictions. The resulting mechanical properties data will be introduced into a more comprehensive and detailed FE model of the structure. Accurate results for static stiffness and detailed deformation patterns for gravity-induced sag and for asymmetric loadings will be obtained from this model. Also, an in-depth study of the impact of kinematic mounting will be undertaken. This will include the interaction between the mounts and the frame as well as local gravity-induced deformations and nonkinematic behavior. The demountable joint scheme will be advanced and a selection of joint materials and radiation/butane resistant, debondable adhesives finalized.

Central region

The design of this region is driven by the goals of providing ultrastable support and services for the detectors with minimum mass. Furthermore, the design must be modular to allow access, and to achieve minimal construction costs. To achieve these goals, the detector layers are paired through common cooling rings to form superlayers. Two superlayers are again paired via a common support cylinder which runs the full 60-cm length of the central region. The individual superlayers are themselves made of three separate detector module units, two of which are 24 cm long, and one which is 12 cm long. These dimensions are fixed by other system constraints.

Thin, 4-mm-high G/C-E ribs are attached to the two long edges of each detector module to greatly reduce gravitational sag of the module, to prevent propagation of any cracks already present in the silicon

wafers, and to provide a robust unit which can be handled in the subsequent assembly stages. The modules are then bonded to the inner and outer surfaces of a pair of cooling rings (one at each end) using a removable adhesive. The result is a rigid, rugged, fully 3-dimensional structure.

The cooling rings are compression molded and constructed of a graphite fiber/cyanate ester material that has been formulated to match the CTE of silicon, and to possess a very small CME. At the same time the material is very low-Z, has low density, and has very high structural strength. The compression molding allows one to precisely and reproducibly match the complicated serrated shape that the cooling rings must have to hold the detectors at the required 7.4° cant angle. Although in many ways the serrated shape is a complication, it does facilitate the achievement of the required circumferential overlap of the detector modules, and at the same time maintains an open structure that allows free circulation of gases inside the STS central region.

The wall thickness of the cooling rings should be as thin as possible in order to minimize both the mass and the path-length between the heat source (the electronics), and the surface which is actively cooled (inner walls of the cooling ring). These walls are a major contributor to the total resistance to the heat flow. From practical considerations, the wall thickness must be sufficient to provide adequate stiffness and to allow manufacture in a consistent, reliable, cost-effective way. The wall thickness is presently $450\text{ }\mu\text{m}$ and is largely determined by manufacturing process limitations.

The use of a continuous cooling ring with a continuous internal wick and fully open gas system, keeps the number of cooling system connections needed to an absolute minimum. That reduces the cost and complexity of the system, and increases the system reliability. The cooling ring also provides a very convenient surface on which to mount the electrical bus which supplies the detector bias voltage and provides the power and the control signals to the chips that make up the electronics for the central region. The rings will also support the optical fibers that carry the data and clock signals. During construction, one has full access to the surface of the cooling ring on which the bus and fibers would run, which will facilitate the assembly of the system. Furthermore, as the cooling rings are already required to be as close to the electronics as possible, only very short jumpers will be required to make the actual connections to the chips.

Each superlayer requires six identical cooling rings (five active cooling rings, and one which serves only as a support ring for the nonelectronics end of the one 12 cm shell assembly). As there are only four different superlayers, only four molds and a total of 24 cooling rings are required for the entire central region. By keeping their number small, and making a pair of cooling rings provide all the services to any complete shell assembly, we minimize the mass in the system and provide a modular system which will allow parallel construction.

To qualify the design, we have used finite element modeling techniques to study the gravitational sag, and possible thermally induced distortions of the module assemblies. The model predictions have been confirmed by measurements. Models of complete shells have also been constructed and the results of gravitational deflection studies are discussed in Ref. 3. The results show that deflections are small and fully compatible with our alignment goals. The models are presently being expanded to study the effects of slight nonuniformity in the cooling ring CTE, and asymmetric heat loads.

Each shell assembly will be supported by one of the two structural cylinders, using kinematic mounts, configured in a 3-point support arrangement. The mounting method will allow the repeatable, high precision location of the shell assemblies, while at the same time maintaining their orientation and the location of their centers, as the entire assembly is cooled down to operating temperature.

Much of the knowledge gained in studying the central region can be directly applied to the forward region. Nevertheless, the forward region has many unique features which are discussed in the following section.

Forward region

The performance of the forward region has the same overall requirements as those for the central region. The same basic concepts that were used in the central region are repeated: cooling rings that double as structural rings and routing paths for the cables, modular construction, and kinematic mounts to support the various components. Similar to the central region, each detector module has G/C-E ribs running along its long edges.

To accommodate the required radial coverage, 8 slightly different wafer types are required to build the 4 different types of modules. The trapezoidal shape of the modules and the radial orientation of the individual strips imply that the strip pitch increases with radius. As the maximum strip pitch is still approximately 50 μm , all the electronics must be near the outermost edge of the modules. Thus, one cooling ring is located near the outer edge of the disk array, and one is located near the mean radius of the array. Finally, the possible slight CTE mismatches between the different cooling rings and the silicon imply that the disk system would be overconstrained if the two cooling rings were directly coupled by the silicon. Thus, two physically separated sets of 12-cm-long detector modules, each coupled to its own cooling ring are required. The two cooling rings will be coupled by three radial webs, located 120° apart in the plane of the disk (see Fig. 4-39). The inner ends of the two 12-cm-long detector module sets will be constrained by flexure mounts, which prevent motion along the beam axis, but do permit radial movement in the plane of the disk. The three radial webs will be attached via kinematic mounts to the space frame.

By having adjacent detectors on alternating faces of the cooling ring, circumferential overlap of the modules is easily achieved. A benefit of the alternating face mounting scheme is that the cooling rings do not have the complicated serrated geometry that they have in the central region, but are just simple rings. As such, the molds to build them are simple and inexpensive. Furthermore, the cooling rings can be produced using a layup technique instead of the much more expensive compression molding technique. In all other respects, the cooling rings for the forward regions are identical to those of the central region, just as the local heat loads are identical.

The cabling plant (fibers and metal) again runs on the outermost surface of the cooling rings. The wide spacing between detector modules on a given face of a cooling ring facilitates the connection of the individual electronics chips to the bus. The active width of the individual detector modules at the radius of the electronics is 6.4 cm, twice the value for the central region. As such the same electronics package can be used for both the central and forward regions, but each forward detector module requires two such packages per detector side as opposed to one for the central region.

The thin planar arrays are naturally much less rigid than the long silicon shell assemblies of the central region. FE analyses were performed to determine the static stiffness and dynamic characteristics of the design developed for this region. This configuration was found to be reasonably stiff, exhibiting a maximum of gravity-induced sag of only 15 μm and possessing a fundamental natural vibration frequency near 70 Hz. The model is presently being expanded to explore other effects such as nonuniform temperature distributions and possible variations of the CTE around the circumference of the cooling rings.

Main power distribution cabling network

Configuring a power distribution system for the silicon tracking system requires a delicate balance between many diverse issues such as:

- a) facilitating assembly, disassembly, and maintenance,
- b) maintaining high reliability,
- c) minimizing magnetic forces,
- d) avoiding induced voltages during magnetic quench,
- e) minimizing shell penetrations, and
- f) avoiding induced loads on the delicate silicon structures.

Table 4-18 lists the combined analog and digital power that must be supplied to the central and forward silicon detector modules. Although the design goal for the electronics is 1 mW/channel, as used in Table 4-18, a value of 2 mW/channel has been used in sizing the cooling system to provide a conservative design margin.

Table 4-18
Baseline voltage specifications.

Main power bus/distribution cables	
Analog chip	3.5 V-1300 A
Digital chip	5.0 V-390 A
Detector bias (maximum)	250 V-6.50 A
Bias voltage lines	1 at +125 V
	1 at -125 V

To facilitate detector assembly, testing, disassembly, and repair of the silicon system, it is clearly advantageous for the individual silicon substructures discussed earlier to be basic units with their own independent set of cables. To supply the 68 cooling rings contained in the forward and central regions, there are a total of 320 cables (160 per end). Thus each cable services typically 20% of the detectors attached to a cooling ring.

To improve reliability, one would like as few connections as practical. However, ease of assembly and repair require strategic placement of numerous connections. To minimize connector mass we have chosen to implement a bonded connection, *e.g.*, "indium scrub bonds," or a "solder" joint. Preliminary tests of a solder joint with beryllium cable material have been completed. This will be an area of continued development in the proposed technical effort.

To supply the large currents needed by the electronics at the large distances (~ 3 m) requires a substantial conductor cross section (~ 20 cm² cross section per detector half). To minimize the pathlength that the conductors present to the particles, the conductors are spread out to as large a diameter as practicable. With respect to Fig. 4-43, the conductors which provide power to the individual cooling rings are brought in along the space frame. The cables placed at this outermost radius result in a nominal thickness of almost 0.6 mm. Beryllium has been chosen as the baseline conductor material to minimize the number of radiation lengths.

Cables will be constructed by laminating flat sheets of specially processed beryllium conductor material with Kapton. Vacuum processing of the beryllium is performed to remove essentially all the BeO from the base material, thus greatly enhancing its ductility. Thin laminated prototype conductors (2 cm wide by 0.5 mm) have been prepared, and proved to be quite formable without evidence of damage.

It would be undesirable to have individual penetrations of the STS outer enclosure shell for each of the conductors at each end. For this reason, the individual ribbon cables join a laminated flat conductor/insulator plate that is sandwiched in between the outer shell and the end closure plate. This laminated plate has embedded in it all the individual power conductors. It is practical to embed the fiber optic ribbon cables in this laminated plate as well. The embedded fiber optic ribbon cables would be potted with an easily removable adhesive in order to permit disassembly.




FIG. 4-43. STS power cabling layout depicted on an arc section.

4.6.2. Silicon electronics cooling

Design requirements

The mechanical stability requirements dictate precise thermal control of the entire silicon array to a nominal operating temperature with minimal thermal gradients. A 0°C detector operating temperature has been selected [3] to keep the detector leakage currents within tolerable levels.

Cooling system description

The requirements of zero vibration, isothermal operation for mechanical stability, low mass, and large thermal loads have led to the development of a liquid-vapor phase change system to cool the detector array. Butane has been selected as the phase change medium based on its normal boiling point (-0.5°C), availability, and cost. Isobutane with a -11°C normal boiling point is also being considered if the temperature drop across the silicon and electronics proves excessive. The cooling system consists of an evaporator located at each electronically heated surface, vapor return lines, a glycol cooled condenser, and a network of liquid feed lines. Figure 4-44 shows a schematic representation of the cooling system design. Subcooled liquid flows from the condenser to the detector by gravity through insulated liquid feed lines. The amount of subcooling is dictated by the heat leaked to the liquid during its transfer from the condenser to the liquid distribution system inside the detector. Once inside the detector the liquid is distributed to the heated surfaces by a combination of gravitational forces and capillary forces. The heat produced in the electronics is transferred to the liquid which evaporates. The vapor is removed from the heated surfaces by pressure gradients generated by the condensing process. The entire phase change cycle is fully contained in an hermetic enclosure which encloses the entire detector array.

Evaporative cooling offers many distinct advantages over other potential cooling methods. It is characterized by high local heat transfer rates. This promotes isothermalization over a large structure. Very low liquid inventories are required in the detector. This minimizes the radiation length contribution of the cooling system. The evaporative cooling system is completely passive. Liquid priming and temperature stability are both self regulating. No pumps or compressors are required during operation.




FIG. 4-44. STS cooling system design schematic.

Wick development

The liquid distribution wick is the key system component in terms of technology development for this application. Liquid for the wick flows from the condenser due to gravity through liquid feed lines. Fine size pores of less than 10- μm -diameter standoff the hydrostatic head between the condenser and the cooling ring and prevent liquid from bleeding from the cooling ring structure. The wick thickness required to distribute an adequate supply of liquid over the heated surface is a function of wick pore size, wick permeability, capillary pumping distance, and power density. The wick pore size defines the available maximum capillary pumping pressure.

The evaporative cooling system uses polystyrene microcellular foam technology that was originally developed for use in inertial confinement fusion systems. These foams have densities of less than 0.1 g/cm³ and high porosity (typically 95%). The pore radius of these samples, determined from scanning electron photomicrographs, varies between 1 and 6 μm . This pore size is acceptable for use as a wick material on the cooling ring. A molding process is being developed for wick integration into the detector cooling rings. The polystyrene foam with its excellent wicking properties, high abrasion resistance, and long radiation length is a highly desirable material for this application.

Safety

The evaporative cooling system is being designed in accordance with SSC specific policies and procedures. These policies incorporate numerous federal regulation, design codes, and orders including: NFPA 101 Life Safety Code, OSHA 29 CFR 1910, DOE Orders 5480.4, and DOE order 5483.1A. The baseline detector design operates with butane vapor at a slightly positive gage pressure. A positive pressure ensures that any leaks flow out of the detector and prevents oxygen and other noncondensables from entering the detector volume. A positive operating pressure requires operation at above -0.5°C for butane or -11°C for the isobutane alternative. In addition, all external liquid and vapor lines are doubly enclosed. The volume around the detector enclosure is continually flushed with dry nitrogen.

4.6.3. Silicon alignment strategy

The STS alignment strategy is to adapt and customize commercially available equipment to achieve the detector alignment requirements described earlier in Table 4-7 and to develop methods for the more difficult task of verification and alignment maintenance. The resolutions of off-the-shelf measurement and alignment equipment are listed in Table 4-19.

Table 4-19
Resolutions of commercially available alignment equipment.

Variable	Resolution
Angle	
Alignment telescopes	15 μ rad
Precision theodolites	3 μ rad
Precision electronic levels	3 μ rad
Precision electronic autocollimators	0.1 μ rad
Length	
Hewlett Packard interferometer	0.05 μ m
Linear Moiré scales	0.50 μ m
Position	
Cameras and centroid software	0.1 μ m
Visual instruments	1 μ m
Cordex measuring machine	2 μ m

The major alignment tasks are:

- a) place the silicon modules onto the cooling rings to form the silicon shell assemblies,
- b) assemble silicon shell assemblies onto G/C-E support shells,
- c) place units from step b) into the space frame,
- d) build the forward region subassemblies and attach them to the space frame,
- e) align the assembled detector with the SSC beam line, and
- f) monitor the alignment status throughout the detector lifetime.

Assembly plan

Construction of the detector's central and forward regions requires placement of about 2700 silicon detector modules to very tight tolerances. To meet the placement requirements an operator-controlled work station with an integrated optical alignment system is used. Special fixtures are used in conjunction with the optical alignment system for the placement of the kinematic mounts which have placement tolerances similar to the detector modules.

The initial step in the construction of a typical shell is the placement and alignment of each cooling ring on a circular plate. This plate is equipped with precision alignment devices for positioning the ring and the kinematic mounts. The optimum center of each ring is determined with an optical alignment system. Each plate, with the cooling ring in place, is then mounted on a cylindrical fixture attached to a single-circle goniometer.

Individual silicon modules are manipulated by the work station. This device consists of an end effector and arm attached to a rotary stage, and three stacked translation stages. Each translation stage is servo driven and includes a linear incremental encoder capable of sub-micron resolution. The end effector is a vacuum platen mounted on a triple-axis flexure mechanism. Each axis is driven by a piezoelectric translator capable of producing angular displacements in the range of 6-8 milliradians. In a typical shell assembly operation a detector module is first retrieved. Part way through the initial motion the module is stopped and adhesive is applied in precisely measured amounts to its ends. Motion of the module then continues to a location adjacent to the mounting facets on the cooling rings. At this location the alignment of the module with the optical reference is determined by viewing fiducials near the ends of the module with a pair of alignment telescopes. If it is misaligned, the module is realigned with the translation stages and piezoelectric translators on the end effector. A single translatory motion then places the detector modules on the cooling rings.

The system described above is used to place silicon modules on the inside and outside of cooling ring assemblies in the construction of central region shells. It is assumed that all the modules are placed on one diameter before the system components are relocated to repeat the operation for the second diameter. For the central region shells, the relocation of system components is simply the translation of the pick-and-place machine, the automatic feed, and adhesive dispensing equipment. All functions of the assembly system are controlled by a multiple axis positioning controller similar to those used in industry for managing entire manufacturing cells.

The kinematic silicon shell mounting concept centers the assembly relative to the composite shell axis, and at the same time restrains the unit from rotation and axial motion. To achieve the desired precise alignment, elements of the kinematic mount that form part of the silicon shell assembly are bonded in place before the shell is removed from the alignment/assembly fixture. As a final step, target fiducials are placed on the silicon shell assembly to mark the optical reference axis to aid in subsequent alignment tasks.

Initial placement and alignment of the two shell support cylinders, with respect to the space frame reference axis, is achieved using simple optical and mechanical measurement methods. During this phase, the shell support cylinders (without silicon shell modules) and mounts are assembled and precisely located.

To achieve the stringent assembly tolerances, a final positioning of the fully assembled composite shell structures must be accomplished. Circumferential alignment of the composite sandwich/silicon shell modules assembly to the reference axis will be achieved by the same technique used to place the individual silicon detector units. Telescopes will be used to align the fiducial markings on the detector units to achieve the 25- μm circumferential placement tolerance. Provision will be made in the design of the kinematic mounts to allow for this fine motion. Once the subassemblies are in place, the individual mounts will be secured in place. Precision movement will be accomplished through piezoelectric linear motion devices which are removed after the alignment is complete.

The method for constructing and aligning the forward region detector units is similar to the technique presented for the central region silicon shell module. The operator-controlled work station will have special adaptive elements for holding and bonding the trapezoidal-shaped detector units. Fiducials on the detector units will be used to locate the elements with respect to the global frame reference on the cooling ring structures.

Individual forward region flat planar silicon structures must be precisely centered and circumferentially oriented with respect to the central region. Location of each planar structure in z must be held to a much less restrictive tolerance. To accomplish these alignment steps, temporary targets will be installed on the large planar arrays while they are being constructed. These targets will be aligned to the global reference frame used in constructing the flat planar array. The targets are used to define the center and azimuthal orientation. The flat planar array center point can be defined by precision placement of cross wires. These targets should be sufficient for placing the unit after the silicon wafers are installed.

The planar array will be set normal to the STS axis by autocollimating off the silicon detector surface. Longitudinal placement will be achieved through an interferometric distance measurement referenced to

the central region origin. Circumferential orientation can be achieved by optically sighting on the targets placed on the cooling rings during the initial assembly.

Detector monitoring

One of the most important tasks is alignment monitoring of the calibrated detector. With a recently developed fiber optic system it will be possible to continuously monitor the relative 3-D position of one subassembly relative to another with a compact, nonintrusive system. For this to be possible the integrity and stability of the subassemblies must be assured so that the subassemblies can be treated as rigid bodies.

The monitoring system is composed of optical fibers installed to monitor the relative positions during assembly, shipping, and between calibration runs. This will be used to keep track of the changes in the carefully assembled parts during these operations. These fiber systems, based on monitoring speckle patterns, will be attached to the relevant subassemblies and routed along the space frame. The light source and measuring electronics required for this monitoring scheme will be outside the detector, leaving no active parts inside the detector.

Alignment and alignment verification are not the only concerns in achieving and demonstrating system stability. Equal emphasis is being placed on component and subassembly testing to demonstrate system stability. Optical techniques are essential to this effort and will be integrated with the overall alignment effort. A relatively new technique, TV holography, will be used to gather quantitative as well as qualitative information about subassemblies and completed detector assemblies with micron resolution over a large field of view. The technique will be adapted to give detailed information about structural integrity, shape, shape changes, structural stability, vibration resonant frequencies and their resonance Q . The results will be used to optimize subassembly designs, aid in the choice of appropriate adhesives, benchmark FE models and provide experimentally tested analytical models for predicting many other detector mechanical performance features.




FIG. 4-45. Completely assembled central tracker

4.6.4. OTD mechanical design

The OTD detector elements that do the actual particle tracking require a support structure to support or fix them at the proper location inside the overall detector. These detector elements or modules described in this report in section 4.4.3 are somewhat robust in and of themselves, but they still require a relatively substantial support structure to assure alignment over their full length.

Developing a sound concept for this support structure is not a trivial matter considering the huge size to weight ratio of the tracker. The 3.4 meter diameter and 8.6 meter length generate a 90.0 cubic meter volume which is estimated to weigh approximately 1800 kilograms. This results in a very light average density of 20 kilograms per cubic meter. The support system must be constructed using minimum quantities of radiation-hard long-radiation-length material while maintaining maximum rigidity and stability. With these goals and restrictions in mind, a design concept has been developed for the support structure for the straw tube modules. This support structure concept will be referred to as the "space frame support system" in this report. Figure 4-45 shows a view of the completely assembled central tracker including modules, and Fig. 4-46 a completed space frame support. A considerable amount of physics evaluation and engineering analysis has been performed on this maturing concept. The methodology dictates that, to achieve our alignment and long term stability goals, the following are required:

- a) An absolute minimum number of mechanical joints between graphite composite components,
- b) Minimization of the tolerance buildup by using only a small number of large components with simple geometry which are thus relatively easy to manufacture to high tolerances, and
- c) Performance of final alignment testing on the assembled support structure in its fully simulated, in-service gravity-loaded support environment.




FIG. 4-46. Completed space frame support system

This strategy reduces costs by requiring precision fits only at the major component interfaces and the module-structure interfaces. Using the space frame support system allows module and structure component fabrication to occur in parallel. This is a major advantage because of its potential to reduce fabrication cycle time and shorten the schedule. There are only two basic types of large components used in the space frame support system. They are large graphite composite cylinders and space frames. The space frame is shown in Fig. 4-46. A support cylinder is shown in Fig. 4-47 and the assembly of the two can be seen in Fig. 4-45. In the following sections, the details of major component fabrications and designs are developed in more detail.




FIG. 4-47. The basic support cylinder for the detector modules.

OTD materials requirements

From the structural engineering standpoint, the material requirements of the central tracker are very restrictive. The support structure must be fabricated with minimum quantities of material but the final structure must have maximum rigidity or stiffness and resistance to deterioration from radiation.

A graphite fiber resin matrix composite has been selected as the basic structural material for a variety of reasons. This radiation-hard material exhibits a high stiffness to weight ratio with a large effective radiation length of 25 cm. The basic fabrication technology exists for the proposed construction of large cylinders utilizing foam cores and large single-unit space frames.

A quasi-isotropic layup has been chosen as the prime construction for all carbon graphite composite components. This layup has six layers with the fibers oriented at 0° , $+60^\circ$, -60° , -60° , $+60^\circ$, 0° to give balanced and symmetrical mechanical properties. Use of the same layup for both the space frame struts and the composite cylinders avoids an unacceptable unbalanced difference of thermal expansion coefficients between these structures. The response of all assemblies to loads including gravity, thermal, and moisture can be expected to be identical. The layup chosen also provides the substantial shear stiffness needed by the cylinders.

The resin systems being investigated include the epoxies and cyanate esters. Cyanate esters are presently the leading candidate resin material because of low moisture absorption and thus smaller changes in elongation for changes in humidity level.

Space frames

The heart of the support system is the space frame shown in Fig. 4-46. The composite space frames perform three basic functions:

- a) furnish four load points that support the tracker on the surrounding detector at the outside diameter;
- b) support the silicon detector at the inside diameter;
- c) register the five superlayer composite cylinders and thus the OTD detector elements themselves.

Each space frame is fabricated from high-stiffness graphite components. These structures are fabricated from struts, joints and rings. The struts, which are the main component in the assembly, are thin-tube-shaped items made by the heat-curing of wrapped B-Staged graphite cloth. The cross sectional sizes of

these struts are 2 cm by 4 cm or 4 cm by 8 cm with a 0.2 cm wall thickness. The joints which connect the struts together are also thin-walled items made by the heat-curing of wrapped B-Staged graphite cloth.

The two sets of five rings which are the interface components to the cylinders are fabricated by hand layups of autoclave heat-cured B-Staged graphite cloth. The rings are manufactured with oversized thickness for later machining by either specifying a thicker cross section or incorporating pads into their thickness.

Each space frame is built by adhesively joining the three components, struts, joints, and rings into an assembly on a large fabrication tool. This tool is not required to have high precision. The tool must be very stiff and fit snugly while rigidly holding the assembled space frame. Once assembled, the tool with the space frame is transferred to a three-axis boring mill on which each of the five oversized rings is machined to a predetermined cylinder matching diameter. The total operation consists of a single setup with only simple two axis moves. Diameter tolerance of 75 μm including a 50 μm concentricity should be achievable.

Support cylinders

These composite cylinders are used as structural members, with an equally important function of furnishing stable bases for the modules. The design features of the cylinders are shown in Fig. 4-47. The cylinder construction uses two composite face sheets and a foam filled core. The most restrictive design requirement for the cylinders is the need to use minimum quantities of material in order to achieve the least thickness in radiation lengths.

A balanced and symmetrical layup with the needed in-plane shear modulus requires 6 plies of unidirectional fiber resulting in face sheets 0.127 mm thick. This is considered to be near the minimum thickness achievable, though the use of Hercules Ultra High Modulus Fiber could perhaps yield two-thirds of this thickness.

Foam cores are needed to give the cylinders some out-of-plane bending stiffness. The core acts as a spacer between the two composite face sheets. To achieve the maximum cylinder stiffness within acceptable radiation length limits, a core thickness of 6 mm is used.

The end closeout design features of the cylinders can be viewed in Fig. 4-47. This area forms the mechanical connection with the space frame. A dense core material is required to withstand the forces that the mechanical connections will apply. This dense core material can range from a denser 300-grade Rohacell foam to graphite composite layups. If a layup is used, the coefficient of thermal expansion must be matched to the cylinder hoop expansion. Although more engineering is needed, these issues do not appear difficult to solve.

The graphite hardware fasteners shown in the Fig. 4-47 are commercially available. Preliminary evaluation indicates that this hardware satisfies the requirements imposed upon it by the interface connection.

The cylinders are to be formed on individual mandrels. These mandrels would be placed in a taping machine, as shown in Fig. 4-48, and 6 layers of the B-Staged lamina would be applied. The foam core which would have previously been thermally formed to the correct radius would be placed on the mandrel and filament applied to hold it in place. After the end closeouts were installed, the last 6 layers of the B-Staged lamina would be applied. Vacuum bag technology would most probably be used to consolidate the composite. The number of cure cycles required, and whether an oven or heater strips would be used to cure the composite, have not been determined. Additional detailed engineering work is needed in this area.




FIG. 4-48. The mandrel used for construction of the cylinders showing the taping operation.




FIG. 4-49. A completed cylinder with shim rings attached.

Shim ring module supports

Sufficiently predictable composite cylinder diameters are impossible to achieve without making several cylinders of each size and iteratively sizing the tooling. Such trial and error sizing would be prohibitively expensive. From the radiation length standpoint, the machining or grinding of a thicker outer wall composite layup is also not an attractive solution. A quality cylinder can be produced on the first attempt by rotating the mandrel during the cure cycle, but the resulting cylinder diameter and outer surface quality will be unknown.

Cylinder shim rings solve the problem of unpredictable diameter and surface quality, as well as other problems. In this scheme the composite cylinders are fabricated in a normal manner. Prior to removing the cylinder from the mandrel, the outer surface is fitted with relatively thick strips or bands of long radiation length material bonded to the composite cylinder. A fabrication such as the one shown in Fig. 4-49 can be made in this manner. These shim rings can then be machined or ground to a precision diameter.

The shim ring machining is done on a two-axis machine equipped with a precision laser encoder that allows 2 to 3 arc-second indexing of the cylinder. This accuracy translates into $23\text{ }\mu\text{m}$ and $10\text{ }\mu\text{m}$ circumferential positioning accuracies for the largest and smallest cylinders, respectively. Linear way systems are commercially available that will carry the spindle and are straight to within 25 microns. It is a simple matter to optically align the mandrel axis to the spindle way system axis to a high tolerance. A $50\text{ }\mu\text{m}$ parallel placement should be achievable. Thus, shim rings can furnish a surface that can be machined or ground to a predetermined diameter and shaped to accept the modules. For the stereo modules this requires an "hour glass" shape.

The detailed view of stereo and axial shim rings is shown in Fig. 4-50. This same view is again shown with the modules in place in Fig. 4-51. Axial modules require simple facets to be machined into the shim rings. This operation requires the spindle way system and the mandrel axis to be parallel and the mandrel to be indexed. Stereo modules require steps or shelves to be machined into the shim rings at approximately 3° intervals. Stereo preparation of the shim rings requires an adjustment in the machine tool such that the spindle way system and the mandrel axis are parallel in the vertical plane but skewed in the horizontal plane by 3° . Very high accuracy of the 3° specification is not a requirement but consistency from module to module is. This consistency is automatically obtained because 360° of modules are done in a single setup.




FIG. 4-50. Detail of the stereo and axial shim rings.




FIG. 4-51. Detail of shim rings with modules installed.

Module attachments

The modules will be located and attached to the shim rings at approximately 80 cm intervals along their length. The shim rings will have fiducials placed in them during the machining operation. These shim ring fiducials will mate with matching module fiducials placed on the module shells during shell fabrication. The module fiducial placement, since it is located by shell tooling, should be repeatable from module to module. The shim ring machined locations should also be very precise since each cylinder will be machined on a single setup and within the accuracy of the machine tool. For purposes of repair and maintenance, the modules should be removable and replaceable from a fully assembled support structure. To allow this, the modules will be installed with removability features to the shim ring cylinders after the space frame support structure is assembled.

An attachment method that has the required features is shown in Fig. 4-52. Provisions for shim ring module attachments are placed in the shim rings during the machining operation that locates the shim ring fiducials. These shim ring attachments will mate with the matching module attachment placed on the module shells during shell fabrication. The concept utilizes an expanding tool that can reach into the inner part of the detector and release the springs that hold that particular module so that it can be removed or replaced. This engineering effort on module attachments is not yet complete.




FIG. 4-52. Module to shim ring attachment.

4.6.5. OTD alignment strategy

There are three levels of alignment in constructing the OTD: the forming of the modules, the module attachment to the cylinder, and the cylinder placement in the space frame. These are discussed below.

Modules

The positions of the straws in the module are determined by the close packing of 159 straws in the stereo modules and 212 straws in the trigger modules. The alignment is set during the formation of the straw bundle before it is bonded into the shell. The outer dimensions are set by a set of clamps that force each layer of straws into a close-packed array at each of the wire support positions. The hard plastic wire support compresses to an accurate position with a precision of about 25 μm . This has been verified by optical measurement of the wire supports and by x-ray measurements of prototype chambers.

Module attachment to cylinder

The straw positions will be determined with respect to the module attachments and also with respect to optical targets which are visible on the outside of the module by sighting along the module. Each straw bundle will be considered as a unit. The position of this unit with respect to the optical targets and the module attachments will be determined by fitting a close packed array to the measured positions of the wires at each of the wire support points. These measured positions can be obtained by examining each module with x-rays. The tungsten wires plus reference elements in each target and module attachment can be clearly seen and measured on an x-ray image. An x-ray examination will be part of the quality check of each module at the assembly point.

The module attachment will have a reference surface at each 80 cm wire support point. The module attachment will be registered to the mating reference surface of the shim ring attachment at each of these points to establish the straightness of the module. In addition, the straightness will be checked using the optical targets.

The shim ring reference surface will be established at each attachment point at the time of machining. This must be done to a mechanical tolerance of about $\pm 50 \mu\text{m}$ at each shim ring. This tolerance is established by the machining operation and will be set and monitored by a laser alignment system on the mandrel milling machine.

Cylinder placement

At the time of the shim ring machining, a set of reference targets will also be set in the cylinder. These will be used during the assembly of the cylinders on the space frame to establish the placement of each cylinder. This alignment can be done optically using standard surveying equipment with a tolerance of $\pm 25 \mu\text{m}$. The space frame ends will have optical targets at radial positions corresponding to each cylinder location. These targets need to be positioned so that they can be viewed from both ends of the tracker in order to align the two ends of the space frame during assembly. The same optical targets will also be used to establish a local coordinate system for checking the positions of the modules and verifying their coordinates with respect to the space frame. The space frame coordinate system is determined by these optical targets and their positions with respect to four load points which are used to support the whole OTD assembly. The load points and the optical targets will be positioned on the space frame during its construction, ideally during the support ring machining phase when the cylinder centers and positions are being machined. Since this is a one piece construction no adjustment is necessary.

4.6.6. ITD mechanical design

The basic detector unit of the ITD is a tile, described in section 4.4.4. Each of the 3200 tiles is a self-contained detector element complete with gas enclosure, front end electronics and high voltage connections. While the mechanical design of the ITD is, as yet, only just beginning, many of the concepts are known. The tiles will be mounted on both sides of a support disk, two such disks forming a superlayer. The disks will be mounted on an inner conic section space frame, with only lightweight stays on the outer circumference. The space frame will be mounted from the end cap calorimeter. Utilities for the first two superlayers of the ITD will be routed along this inner conic section space frame, behind the ITD, and out through the notches between the end cap and barrel calorimeters. Utilities from the third superlayer ($z = 4.26$ m) will run radially out into the notches between the barrel and end cap calorimeters.

Microstrip tiles

The tiles, after passing an exhaustive test procedure, will be accurately placed on the support disk. Fiducial marks on the tiles and support disks will be used to ensure correct placement within prescribed tolerance. The design of the tile, and the tiling onto the support disk, will be arranged to optimize the (r, ϕ) coverage and minimize cracks in the layer. The two radial layers and the two stereo layers of each superlayer will be staggered in (r, ϕ) to ensure that there is complete coverage in the superlayer. The tiles will be overlapped in theta, to obtain complete coverage of each eta bin, as shown in Fig. 4-53. The tiles will be mounted on the support disk using dowel pins, or a similar mechanism, to facilitate replacement of complete tile units during long shut-down periods.




FIG. 4-53. Tile mounting to support disk.

Cooling

The electronics lie on rings, approximately 2.5 cm wide and with one ring per radial segmentation of the microstrips. The bulk of the heating in the detector will be concentrated in these rings, each of which will dissipate 200–400 W (20 to 40 tiles). Cooling rings are therefore required as an integral part of the mechanical support structure and are mechanically connected to the electronics. The design of these rings is not yet settled, and it is not clear whether forced “air,” liquid or evaporative cooling is required. Whatever system is chosen, the mass of that system will be an important consideration.

Support disks

The support disks will be made from carbon fiber laminated on a foam core to provide a rigid lightweight structure of low Z material. Close contact will be maintained with the design of the OTD cylinders and silicon tracker, to avoid duplication of effort and facilitate the design. The low voltage distribution to the front end electronics will be an integral part of the design, and may well replace some of the carbon fiber in the laminate. The disks will be connected together using lightweight spacer rings to form a superlayer. Cut-outs in these rings will allow passage of cables and utilities.

The space frame

The superlayers will be mounted from a conic section space frame as shown in Fig. 4-54. Detailed design of this structure will be based on the principles and materials used for the OTD and silicon tracker space frames. The third superlayer will have an extra stiff disk between it and the end cap calorimeter. This small amount of material will be out of the tracking volume, hence have minimal effect on the physics, but will add greatly to the rigidity of the ITD.




FIG. 4-54. Support cone to which support disks are attached to provide mechanical support and positioning of the tile planes.

The ITD will be mounted from the end cap calorimeter via the inner space frame and extra stiff disk behind the third superlayer. Mounting the ITD off the end cap calorimeter has many advantages. Access

to the ITD, OTD and silicon tracker can be obtained without disconnecting cables of read-out fibers. The relative internal alignment of the ITD can be done during manufacture, and will not be disturbed during access. This is most important as the internal alignment forms the basis of the Level 1 trigger. Opening and closing the end cap calorimeter will result in some displacement of the ITD relative to the silicon tracker and the interaction point. There are six degrees of freedom of this displacement, three axes of rotation and three of translation. The magnetic forces between the end cap and barrel calorimeters will ensure that the ITD position is reproducible in two of the three axes of rotation, the critical axes for the Level 1 trigger, and in z . The position of the ITD, after opening and closing the end cap, will vary in x , y and ϕ , and then only by small amounts (less than 1 mm in x and y , and 0.1 mrad in ϕ). These errors of alignment will not affect the operation of the Level 1 trigger.

In the tracker layout there is a 20 mm clearance between the ITD and both the OTD and silicon tracker. The alignment of the ITD must be much better than this, hence mechanical interference between the tracking sub-systems will not be a problem during insertion. Remote position sensors will be employed to measure the relative alignment of the tracking sub-systems after insertion, after turning on the magnetic field, and to ensure there is no vibration. These sensors will be monitored during insertion to double check that no mechanical interference occurs.

4.6.7. ITD alignment

There are two alignment requirements for the ITD. The first is the internal placement accuracy of the device and the second is accuracy with which its position is known relative to the other tracking devices. The internal accuracy concerns the placement of the tiles on the disks and the disks on the cones. This affects the Level 1 trigger accuracy, which depends on the measurement of ϕ from layer to layer. The required placement accuracies of each tile relative to the others in the ITD is 40 μm in the circumferential direction and 1000 μm in the radial direction.

The alignment of the ITD and STS affects the momentum resolution of the detector. Since the ITD is mounted from the endcap calorimeter, it will not necessarily return to the same position after withdrawal and reinsertion. It will also move when the magnetic field is switched on. To keep track of these motions, remote position sensors will need to be used to measure the relative ITD–STS location. These will provide the required input data to enable the relevant constants to be updated. The final relative alignment between the ITD and STS will be established using tracks which connect the systems.

4.6.8. OTD and STS supports

In contrast to the ITD which is supported from the endcaps, the STS and OTD are supported from the barrel calorimeter. A highly stable, integrated support structure system that maintains a precise physical relationship between the outer tracker straw-tube array and the silicon tracker strip detectors is essential. Independent support concepts for the outer tracker and silicon tracking system have been defined, but significant engineering design studies are needed to achieve an integrated design.

The final integrated support structure must provide a stress-free stable alignment reference system in the presence of:

- a) environmental changes such as pressure, temperature, and humidity;
- b) external loading, such as shock and vibration;
- c) component material creep.

Our present studies are centering on how to approach these various design parameters with a cost effective, passive support system that provides a reasonable in situ access to the very large central tracking system with all of its complex array of data readout, power cabling, and service connections. Clearly, much work remains to be done, although at this juncture we see no reason for deviating from a passive system.

The Outer Tracker G/C-E composite shells are supported from opposing ends with two space frames. Extensive calculations have been performed to quantify the sag of the straw-tube system under gravity

loads, including the gravity loading effects of the STS. These studies included estimation of dimensional changes due to thermal and moisture effects.

The OTD support structure is very stiff but somewhat fragile. Large differential deflections between the tracker structure and the available mounting surfaces will result from relatively small temperature changes. These deflections are generated by differential expansion between the steel calorimeter with its associated high thermal coefficient versus the graphite tracker structure with its near zero coefficient of expansion. Shown conceptually in Fig. 4-55 is a system which accommodates radial and longitudinal differential expansion in a stress-free manner by allowing controlled relative motion while still anchoring one corner of the tracker and maintaining position alignment. Radial and longitudinal motion will be accommodated by either the use of sliding low friction mounts or tangent bar beam bending support.




FIG. 4-55. Outer tracker kinematic mounts.

The OTD mounts will be adjustable, so that the trackers can be aligned along the magnetic axis of the solenoid, and centered on the beam. Major optical targets at the support points will reference the tracker position with respect to the magnet survey coordinate system.

The STS is supported at the innermost diameter of the OTD conical shaped space frames. Mounting pads, with associated bolt patterns, located on opposing ends of the two space frames, will be constructed so that they are precisely aligned, radially, circumferentially, and longitudinally. This feature will aid in the initial alignment of the STS. However, with time, it is expected that small deviations in this precision set-up will occur, i.e., either due to material dimensional changes from environmental effects or possibly from moving the massive detector end caps. In either case, the STS support concept provides some accommodation for changes in $r(x, y)$ and z . In contrast, the OTD space frame rigidity has been sized to limit these excursions, and frame materials were chosen to avoid significant dimensional changes. Tangential (ϕ) misalignment of the two space frames is not likely because of the manner in which they are supported by the calorimeter.

Initial alignment of the STS with respect to the OTD will be performed without the large silicon array installed inside the STS enclosure. STS weight and balance will be simulated during the alignment procedure, thereby permitting reinsertion of the fully assembled STS without need for further adjustment. Once the STS enclosure is completely aligned with respect to the OTD, the mounting pads would be fitted

with precision pins that insure repeatability of placement. It is important not to expose the silicon system to the prospect of damage during the initial fit-check.




FIG. 4-56. OTD/STS support interface demonstrating the STS "tangent bar" kinematic supports.

A triad of support beams, fabricated from zero thermal expansion and low moisture expansion G/C-E composite, form the STS kinematic support feature, as seen in Fig. 4-56. The beam triad is referred to as a "tangent bar" concept. Mounting pads for the STS beam/shell support points are spaced at equal 120° intervals around the shell perimeter. This arrangement provides a degree of radial flexibility through beam bending, while affording considerable tangential and longitudinal stiffness. Small diametrical changes in the shell, possibly due to slight fluctuations in barometric pressure, will not cause any radial misalignment of the STS with respect to the OTD space frame. The two tangent bar mounts at opposing ends of the STS differ in one small respect. At one end, the beam triad mounting face would incorporate a flexure to accommodate a small amount of longitudinal movement between the STS and OTD. This feature avoids over-constraining the STS in z . Even though expansion effects have been significantly limited, the overall 5-m length between support points could result in finite movements and need for accommodation.

4.6.9. Servicing and utilities

Table 4-20 provides a breakdown of the utilities presently envisaged for the central tracking system. The utilities needed for the OTD and STS will be routed along the space frame struts which support the OTD. They enter the tracking volume through space provided between the barrel and endcap calorimeters. ITD utilities are routed behind the ITD detectors and pass through similar spaces in the calorimeter system. They must be able to move with the ITD when an endcap is retracted. The large numbers of cables and fibers present an important integration problem which will require significant attention as the detector design progresses.

Power supplies for the central tracking subsystems are mounted on the outside calorimeter perimeter. This location is easily accessible and minimizes the length of power distribution cables and associated voltage drop. An abrupt loss of a major portion of the central tracking subsystems can be avoided by segmentation of the power distribution feeding the large array of detectors. A reliability study needs to be conducted to optimize this aspect of the design.

The access required for repair and maintenance will be at two basic levels. Work on the power supplies, signal distribution and gas distribution system can be done by entering the region just outside of the barrel calorimeter. Since this area is easily accessible, the time for quick repairs of these systems is a few hours.

Table 4-20
Utility specifications for the tracking detectors.

Service	Subsystem	Requirement	Specifications
Electrical	OTD	Low voltage for electronics power — 1.8 kW per end	5 V dc, 1 amp per module, total 360 amps per detector end supplied via 360, 5 mm × 0.5 mm laminated Kapton/Copper cables. Total 360 power supplies, 12 per crate.
	OTD	Sense wire high voltage - power — 0.15 kW per end	2000 V dc, 1 microamp/wire, on the average 200 micro-amps/module. Each end is supplied via 360, 8 m long by 3 mm diameter coaxial cables.
	STS	Low voltage for electronics	3.5 V dc, 1300 amps
		and voltage for detectors	5.0 V dc, 390 amps
		power — 3.25 kW per end	250 V dc, 3.25 amps
			Each end is supplied via 160 cables.
Module Gas	ITD	Low voltage	Not yet designed in detail, 3.5 V dc to 5.0 V dc. Estimate about 1000 amps.
	ITD	High voltage	Back plane voltage: 10 per layer = 120 per end Drift plane voltage: 10 per layer = 120 per end Cathode strip voltage: 1 per tile = 1560 per end All voltages ~700 V, 200 μ A per tile Each end supplied by 1800, 10 m long by 2 mm diam. coaxial cables.
	OTD	CF ₄ -Isobutane (80/20%)	Supplied to straws via manifold. Gas flow 2 cc/sec per module, nominally 0.72 liters/sec per detector end. Gas feed and return lines 2.5 cm diam., 10 total.
	ITD		Composition not yet known. Gas supplied to tiles via a manifold on each superlayer. Total gas volume = 50 liters per superlayer. Gas supplied to superlayers by 1 cm diam. feed and return lines, 6 per end. Gas velocity in line = 17 cm/sec.
Detector cooling	OTD	3.6 kW total heat load	Convective cooling of on-board electronics with dry N ₂ gas fed via a manifold system, 80 liter/sec total flow, using 20 1 cm diam. inlet and exit lines per detector end.
	STS	6.5 kW total heat load	Silicon detector electronics cooling with butane evaporative system. Passive thermosiphon arrangement with one liquid and vapor co-axial feed/return line per detector end. Liquid inlet 0.028 liters/sec, 2 cm diam.; Return 6 liters/sec, 5 cm × 9 cm rectangular duct.
	ITD		Not yet designed.

(continued on next page)

Table 4-20 (cont'd)
Utility specifications for the tracking detectors.

Service	Subsystem	Requirement	Specifications
Instrumentation	OTD	Position sensors	Monitor centroidal motion of straw modules relative to silicon to within 15 μm .
	OTD	CF ₄ -Isobutane leak detection	Four monitoring points per superlayer. Sense presence of gas leak into the N ₂ inerting gas space.
	OTD	Temperature sensors	Sense temperature of electronics and N ₂ inerting gas volume. Sensitive to 1°C, placed at 4 points per superlayer.
	OTD	Voltage sensor	Monitor electronics and sense wire voltage conditions, 1% resolution. Location and number dependent upon final design.
	OTD	CF ₄ -Isobutane gas purity	Monitor recirculating CF ₄ -Isobutane gas. Sensitivity: < 10 ppm, HCl, HF, F ₂ , Cl ₂ , Br ₂ < 50 ppm, O ₂ /CO ₂ + CO < 500 ppm, Ar, N ₂ , H ₂ O
	STS	Silicon detector module temperature	Monitor electronics/silicon strip detector temperatures using integrated sensors in the electronics chips.
	STS	Butane vapor space temperature and pressure	Monitor butane vapor pressure and temperature conditions at 4 locations, 1% resolution.
	STS	Silicon detector stability	Fiber optic sensors placed at nominally 100 locations to monitor relative motion of silicon shell and forward planar array module motions. Resolution 1 μm . 50 fiber optic cables/detector end.
	STS	Voltage Sensors	Monitor electronics voltage supply at silicon detector.
	ITD	Position sensors	Monitor centroid motion of ITD relative to silicon to within 15 μm .
		Leak detection	Four monitoring points per superlayer.
		Temperature sensors	Four monitoring points per superlayer. Sensitivity, 1°C.
		Voltage sensors	Monitor electronics and high voltage conditions. Location and number dependent on final design.
		Gas purity	Monitor output gas. Sensitivity: < 10 ppm, HCl, HF, F ₂ , Cl ₂ Br ₂ < 50 ppm, O ₂ /CO ₂ + CO < 100 ppm, Ar, N ₂ , H ₂ O

Access to the devices inside the calorimeter will be much more time consuming. This will require the extraction of the endcap parts of the SDC detector and would allow a 4 m long access space. An estimate of the required time is about 10 days for extraction and reassembly of a forward region.

The most likely need for access in the OTD case is to replace an electronics card. Extensive burn-in time prior to close up will be required to keep the electronics failure rate low. There may need to be access to repair a high-voltage short if a straw tube wire breaks. We are investigating the possibility of using remote fusing to remove individual wires from the high voltage distribution.

In situ maintenance of the STS will be quite limited. Access to the STS outer enclosure ends can be provided by withdrawing the endcap calorimeter. This would provide sufficient clearance to work on the external services of the STS. Problems with electrical and fiber optic connections could then be resolved. Also, should a leak develop in the butane supply or return lines, it would be possible to correct this problem. However, internal problems in the STS can only be resolved by removing the STS during scheduled extended down times. It is estimated that the STS could be removed and completely serviced and reinstalled in 3 months.

Access to the OTD and ITD gas system should be straightforward in the gas mixing areas and in the region outside of the barrel calorimeter. If a leak were detected in the tracking region, it would be necessary to access the front end electronics area, determine the point of the leak, and repair the leak or isolate the module from the gas system.

4.7. Research and development plan

4.7.1. STS R&D and program plan

Detectors and electronics

Development of the detector modules combines R&D on detectors, front- end electronics, high-density electro-mechanical integration, data transmission and cabling, and mechanical design. Novel designs in all of these areas are required to bring the Silicon Tracker to realization. The overall system concept and details of implementation are strongly interdependent and work on both levels is being pursued concurrently to develop a production-ready design in time to initiate mass production in 1994. A well-considered combination of basic measurements and system tests is critical to timely completion of the project.

The behavior of the silicon strip detectors is being studied both through detailed calculations and measurements. Determination of detailed pulse shapes in the presence of a magnetic field is essential for realistic simulations of the electronic design. Changes in detector characteristics after radiation damage must be determined through the lifetime of the detector. This work is well underway and will proceed into 1993 when the detector design will be set.

The electronic design is closely coupled with the results of strip detector R&D. Prototype circuits are needed to test detectors, and vice versa. The design concept of the analog circuit blocks has been determined, detailed circuit implementations have been studied and refined, and a complete analog front-end system submitted for fabrication in early 1992. Various circuit variations for the digital circuitry, *i.e.*, time stamping, data buffering, and on-chip data sparsification have been studied. A preliminary digital readout IC is in preparation to be available in mid-1992. Beyond tests of circuitry and detectors, these ICs will be used to assemble a test detector module with output cable that incorporates key features of the final module design. This module will be capable of operation at the speeds anticipated in detector operation at the SSC, and will be used to test system aspects and develop the final design of the electronics, module arrangement, and cabling scheme. The prototype module will also allow measurements of system performance in the course of radiation damage. Module tests are scheduled to begin in the 3rd quarter of 1992. Results of these tests will determine the final design of the complete detector module and the data transmission scheme. Refinements of the electronics and the module will proceed into 1993, culminating in a preproduction prototype in mid-1993. This final prototype will undergo complete system tests before and after irradiation before mass-production commences in 1994.

The detector module is conceived as a self-contained detector subsystem that can be operated by itself and tested completely prior to installation in the support and cooling structure. This allows extensive quality assurance measurements at all stages of the assembly process.

Mechanical engineering

A conceptual design that addresses *all* aspects of the compact and highly integrated STS must be developed with difficult areas identified and studied in detail at an early stage. The mechanical engineering and design plan is logically organized into a detailed series of prototype developments leading to the final fabrication and assembly of the detector system. The issues that are being addressed by this effort focus on four main topics:

- a) fabricability of ultralightweight composite structures,
- b) detector stability, *i.e.*, long term structural behavior,
- c) cooling a highly dispersed array of electronics, and
- d) survivability of materials and their mutual compatibility in a high radiation environment.

Considerable progress has been made to define the mechanical, cooling, and structural arrangement of the silicon tracking system. Major milestones attained thus far include:

- a) proof-of-principle of the butane evaporative cooling system using prototype polystyrene wicks,
- b) compression molding of 30° arc segments of a cooling ring using AS4 graphite fiber/phenolic resin,
- c) structural prototype of 24 cm long edge bonded silicon wafers reinforced with graphite/epoxy ribs,
- d) 1.0 kg/m² areal density truss core panel and shell construction demonstration,
- e) 5 μ m stability of truss core panel after exposure to 1×10^{15} n/cm², and
- f) butane and graphite/epoxy compatibility verified by tensile specimen results, before and after irradiation.

Nevertheless, many significant issues remain to be addressed before it is prudent to proceed with production of costly major items. These remaining engineering activities form the basis of our R&D and final design efforts.

In constructing the STS program plan, we identified the central region as the first major subsystem element on the critical path. Solving the technical issues embodied in this area would serve to greatly diminish technical risks in many related areas. This is because elements contained in this region are basic building blocks for the entire detector. To this end, steps are now being taken to:

- a) develop the fabrication techniques for the cooling ring, and associated polystyrene wick structure,
- b) construct a full prototype silicon shell with simulated electronic heating,
- c) demonstrate precision construction of silicon detector arrays requiring highly repetitive assembly techniques,
- d) perform prototype silicon shell structural stability tests, and
- e) demonstrate stable operation of the hydrocarbon thermosiphon cooling system.

Concurrent with this effort, we will be conducting essential material investigations and compatibility tests. This effort is phased throughout the above activities in order to ensure the appropriate materials are being used in the prototype demonstration phase.

A prototype graphite support cylinder and associated kinematic mounts, which support the silicon shell assemblies, will be used in the final stability tests. Fabrication of the graphite support cylinder will be a precursor to constructing the very large G/C-E cylinder that must ultimately be constructed for the outer enclosure. Both concepts employ a truss core sandwich structure. Our program plan is based on prototyping wherever practical and thereby lowering the risk as the program progresses. It is reasonable to expect that the prototypical units may actually be used in the final assembly contingent upon serious problems not developing.

Once the primary silicon detector module mounting and cooling issues are solved in the central region, work will commence on the forward region. The flat planar arrays have distinctly different structural aspects that must be prototyped. Nevertheless, the method for cooling the large electronic module array will draw heavily upon the successful demonstration of the central region prototype.

Work on the space frame and outer enclosure will start shortly thereafter. These items are major construction elements that have a fairly long lead time. The metal matrix space frame possessing demountable joints presents a challenge and activities in this area cannot be delayed for too long.

Prototyping of the main power bus and distribution cables will be interwoven amongst the above activities. Laminating the ribbon cables using beryllium is not anticipated to be a serious problem. The actual cable layout, however, and the cable connection scheme do have serious ramifications on the overall detector design. A well thought-out design approach that includes mounting and disassembly of major components is essential before fixing the final design of any of the major components such as the space frame. If not carefully thought out, the cable layout has the potential of making it impractical, if not impossible, to service or maintain the STS.

4.7.2. OTD R&D and program plan

A major achievement of the OTD R&D program has been the construction and assembly of a 4-meter-long stereo module shell prototype that meets the tolerances in rigidity and straightness. Measurement of efficiencies and resolutions using prototypes of the electronics designed for these straw tubes are now in progress. We need next to establish that we can also construct a trigger module shell and align the straws to the required tolerances. The inside surfaces of the trigger module shell are curved so that the wires lie along a radius. The design of the final wire support and the resistive terminator will soon be completed. There are numerous electrical and mechanical details of the endplate connections that must be optimized. We plan to build a total of six stereo modules and six trigger modules, all of the final dimensions, so that we can assemble a multisuperlayer prototype tracking system sector consisting of two axial superlayers and two stereo superlayers, each three modules wide. These will be used to develop the large-scale fabrication techniques for the modules as well as the alignment techniques. During the remainder of 1992, we must establish that we can operate the modules with $\sim 100 \mu\text{m}$ resolution at a gas gain of $2-5 \times 10^4$ using prototypes of the final front end electronics.

The multisuperlayer prototype system will be used to test the principles of noise and oscillation tolerance. This will require the development of cross-talk free and low-noise connections between the wires and the electronics. Once we have established the operating conditions, we can develop the final packaging of the electronics.

Prototypes of the multichannel preamp/shaper/discriminator with tail cancellation will be built and tested on the prototype modules. Prototypes of the time measuring circuitry and of the synchronizer for the track segment trigger will be built and tested on modules. These will allow a full system test in early 1993.

Straw tube R&D requires completion of aging studies to determine that the materials used in the module and straw construction and the electronics are sufficiently radiation-hard, including irradiation with neutrons. We will continue to study chamber lifetime relative to SSC operating conditions, including high-rate studies using prototypes of the final electronics.

Substantial R&D on the support structure remains. We will build and test flat prototypes of the support cylinders and shim rings to verify the mechanical properties. The module attachment structures must also be prototyped. We will develop prototypes of the space frame support struts and flange cylinder connections. The structural properties from the prototypes will be used to verify the finite element analyses needed to complete the OTD support design.

4.7.3. ITD R&D and program plan

It was only in 1988 that Oed [34] built the first GMD and observed the phenomenon of electron multiplication in gases with microstrip anodes which were photolithographed on thin glass substrates. Soon thereafter, the Pisa and the NIKHEF groups [36,37,40] fabricated similar devices on other substrates and demonstrated that GMDs are intrinsically capable of achieving superior energy and spatial resolution and better rate capabilities than conventional wire based devices. The GMD technology has made tremendous progress in the past three years. As an example, the NA12 experiment at CERN uses 8 small GMD planes with a total of 1100 readout channels [38]. Recent work by Bouclier *et al.* [43] attests to the impressive rate capabilities that are, in principle, possible with present day GMDs.

However, much work still needs to be done before GMDs free from time and rate dependent gain instabilities and aging effects can be routinely used in large-scale experiments. Gain stability, which is of primary concern, is intimately linked with the search for a substrate material with suitable resistive properties. Proper electrostatic design of the GMD is also essential. The problems of efficiency and aging are related to the choice of gas, the metalization and the radiation hardness of the components. In addition, there are specific R&D issues related to our present ITD design, such as radial anode geometry with keystone cathode structure and the fabrication of large 20 cm \times 30 cm photolithographic masks and prints.

There are a large number of HEP centers which are working on generic GMD issues. We are thus in the fortunate position that several ITD R&D problems, such as substrates, gases, aging and electronics are being addressed by many institutions throughout the world. Within the SDC, there are working GMDs at Carleton/CRPP, Liverpool, Rutherford-Appleton Laboratory (RAL), and Texas A&M. We describe below the R&D activities planned within the SDC.

Substrates

The right choice of substrate is essential for the stable and rate independent operation of GMDs. Thus there has been considerable R&D activity in this field. Several substrates have been found by workers within and without SDC, which may be usable, but perhaps not ideal. These include Tempax glass, Tedlar and boron ion-implanted quartz plates. There are other encouraging candidate substrates such as ion-implanted Kapton and types of iron oxide glass which are also being investigated and may lead to other solutions.

RAL has been evaluating and will continue to evaluate different glasses as suitable substrates, with emphasis on gain stability, practical limits on the usable thinness, and rate performance. There is ongoing R&D on large area lithography on glass in collaboration with industry. RAL has received substrates with dimensions of 178 mm \times 178 mm.

The Carleton/CRPP group is working on a collaborative R&D project with the Sauli group at CERN on the development of GMDs on plastic substrates. Precision 200 μ m pitch GMDs have recently been produced on white Tedlar with the help of a microelectronic firm in Montreal, and are currently under test.

GMD R&D on silicon dioxide/silicon substrates is being carried out at Liverpool, at the University of Montreal and at Texas A&M. These substrates have the additional advantage that the front end electronics can be fabricated on the substrate material itself. Texas A&M is also investigating orientation dependent etching on silicon substrates to produce chambers with 75 μ m anode to anode pitch.

The Texas A&M, Rochester and Purdue groups have been carrying out a thorough and systematic evaluation of candidate plastic substrates. Several candidate substrates in the correct resistive range have been found. A new material developed by B.F. Goodrich combines standard thermoplastics such as ABS or polystyrene with a conductive copolyether in sufficient concentration that the copolyether forms quasi-continuous strands in the parent plastic matrix. This prevents the tunneling or hopping form of non-ohmic conduction in the substrate. It is also possible to adjust the resistivity of Kapton and other polyamides using ion implantation. These materials are claimed by the manufacturers to be metalizable

and etchable and will be evaluated. Other substrates that are also under investigation are polyurethane, amorphous silicon, and germanium.

Metalization

The metal chosen for anode and cathode strips should have good conductivity and good adherence to the chosen substrate. R&D on metalization is being carried out at all collaborating ITD institutions to find the best metal, or combination of metals, and to optimize the plating thickness. Other factors affecting the choice include aging, rate capability, and gas gain. All these factors are interdependent with the choice of gas mixture.

End termination

Field lines concentrate near the end points of the anode and the cathode strips in a GMD. To reduce the field in these regions and to minimize the probability of breakdown, the anode and the cathode strip ends may be shaped to avoid sharp corners. Passivation of the region is also a possibility. Various termination geometries will be tested in order to minimize the probability of breakdown. The techniques of passivation at the end of anode and cathode strips will also be tested.

Aging and rate dependence

The GMD-ITD has to be able to operate for 10 years at several times the design luminosity. In addition, the detector performance must be constant despite variation in instantaneous rates.

Plastics and most amorphous materials such as glass are probably affected only by the specific energy loss and the damage is independent of ionizing species. One may therefore use charged particles, neutrons or photons to study radiation damage in candidate GMD substrates. However, particularly for plastic polymers, it is insufficient to make before and after exposure studies. It is important to measure parameters such as resistivity during an exposure. Thus, some sources might lend themselves more easily to such a study than others.

The aging characteristics will be tested initially using high intensity x rays followed by beam tests to understand rate effects. Using x rays, tests of up to several Coulombs per cm per strip will be done.

This R&D program on the basic operation of gas microstrip detectors can be carried out using small (10 cm × 10 cm) detectors. Plates of this size can be rapidly fabricated ensuring rapid turn-around of new designs.

Positional accuracy and two-track resolution

While these parameters may be estimated from the geometry of the detector, experimental confirmation will be required. A test beam at RAL, equipped with silicon microstrip detectors and a read out system, will be available within the next few months. This may be used for an indication of the performance, but the energy of the beam is low, 700 MeV protons. High energy beams will eventually be required for definitive measurements.

Timescale

Most of the above extensive R&D program is of a generic nature and will be undertaken in conjunction with groups not allied to SDC. In order to meet the SDC schedule we must build a large scale prototype detector, before proceeding to production, by early 1994. This implies the completion of the R&D program in two years, 1992 and 1993. This is ambitious, but not unrealistic given the worldwide resources being devoted to the basic work.

4.8. Scintillating fiber charged particle tracker

4.8.1. Introduction

The baseline straw-tube tracker described in the previous sections has occupancies which, for the inner layers, are just acceptable at SSC design luminosity. The scintillating fiber technology can potentially provide an outer tracking system of occupancy smaller by about an order of magnitude, shorter response time and better granularity than our baseline straw system. However, there is much less experience with scintillating fiber trackers than with drift chambers, and further R&D is required to provide the technical basis for an informed choice between the two technologies. We therefore consider the scintillating fiber tracker as an option for the SDC OTD and describe here the parameters for such a tracker, the R&D status and the future directions.

Design concept

The detection elements are 925 μm diameter scintillating fibers optically coupled to clear fiber waveguides of 1 mm diameter which "pipe" the light to photodetectors situated outside the tracking volume. The scintillator and waveguide fibers are fabricated from polystyrene. The scintillator contains two fluorescent dyes with light emission in the yellow-green ($\lambda \sim 530 \text{ nm}$) to maximize optical transmission and minimize the effects of radiation damage. The photodetector of choice is the Visible Light Photon Counter (VLPC) developed by Rockwell International Science Center [44], a device which combines high rate capability with high quantum efficiency at visible wavelengths (QE $\sim 85\%$ at $\lambda = 565 \text{ nm}$). Amplified signals from the VLPC's are input to the associated readout electronics and control systems. Trigger and data acquisition functions are to be combined in application specific integrated circuit (ASIC) technology and will include appropriate segment-finding and linking algorithms for Level 1 triggering, as well as pipelined Level 1 and Level 2 data storage.

As shown in Fig. 4-57, a tracking system of $\sim 475 \text{ K}$ fiber channels is proposed, which consists of six superlayers of fibers (B1-B6) arranged in concentric cylinders, covering the region $|\eta| < 2.3$ and $0.60 \text{ m} < r < 1.65 \text{ m}$. The design takes advantage of the low occupancy of the fibers to permit the placement of superlayers at relatively small radii ($r \geq 50 \text{ cm}$) relative to the beam to maximize η coverage for tracking and triggering. The all-cylindrical geometry allows for uniform coverage over the full rapidity interval without cracks or clusters of localized material within the tracking volume, such as support structures, detector components, utilities, or services. Because of the substantial coverage in η , the design obviates the need for, but does not preclude, a separate intermediate tracking system. Table 4-21 lists the characteristics of the superlayers.

For a robust trigger, measurements from at least two out of three superlayers are required to achieve good rejection of fake tracks. Superlayers B1, B2, and the axial fiber layers of B3 are used for triggering and cover $|\eta| \leq 2.3$. Outer layers (such as B4) may be included in the trigger logic to increase redundancy of measurement as the machine luminosity approaches $10^{34} \text{ cm}^{-2}\text{s}^{-1}$.

Good pattern recognition requires relatively uniform spacing of measurements throughout the tracking volume, and good momentum resolution requires measurements over as large a radial lever-arm as possible. Layers B4-B6 provide the required extended radial coverage. Cost considerations and the materials budget (in radiation lengths) limit the number of these superlayers to three.




FIG. 4-57. Schematic of the central tracker exhibiting the fiber superlayers.

Table 4-21
Parameters and characteristics of the superlayers of the scintillating fiber tracker.

Fiber super-layer	Radial location (cm)	Fiber layers x =axial u =stereo v =stereo	Fiber channels per end	Scifi length for axial fibers (m)	$\Delta\eta$ coverage	Wave-guide length (m)	Total fiber length (m)	Expected mean no. of photo-electrons detected for 925 μm diam. (at $\eta = 0$)
B1	60	$2x, 2x$	15.9K	3.00	2.3	6.93	9.93	4.6
B2	76	$2x, 2x$	20.1K	3.65	2.3	6.12	9.77	4.6
B3	92	$2x, 2x$	24.4K					
		$2u, 2v$	24.4K	4.30	2.3	5.31	9.61	4.5
B4	108	$2x, 2x$	28.6K	4.30	2.1	5.15	9.45	4.6
B5	136	$2x, 2x$	36.0K	4.30	1.9	4.87	9.17	4.7
B6	165	$2x, 2x$	43.6K	4.30	1.7	4.58	8.88	4.9
		$2u, 2v$	43.6K					

4.8.2. Mechanical structure

The basic construction element of a superlayer is a doublet ribbon of 1024 fibers, placed onto a precision stable-base cylinder of composite material (see Fig. 4-58.) The composite is selected for strength, rigidity, and low mass, and is made of 0.1 mm thick graphite-epoxy face sheets sandwiching a 1.6 cm thick Rohacell foam core. This combination provides a radiation length of less than 0.35% per cylinder at $\eta = 0$. All superlayers B1-B6 support two doublet ribbons of axial fibers, one on the outer surface of the support cylinder and the other on the inner surface. Additionally, B3 and B6 support four layers of stereo fibers—two layers forming a doublet with $+6^\circ$ stereo and two layers forming a doublet with -6° stereo. The stereo layers are supported on the outer surfaces of the cylinders atop the axial fibers. It should be

noted that the layers of stereo fibers are situated outside the volume used for triggering, to minimize the material at the trigger level.




FIG. 4-58. Schematic of a superlayer having axial fiber doublet ribbons mounted on the inner and outer surfaces of a support cylinder.

The support structure for the stable-base cylinders serves to locate all of the components of the SDC charged particle tracking system, including the silicon tracker and other tracking devices, and must meet the precision alignment requirements for locating the tracking elements specified in Table 4-7.

Figure 4-59 displays the main support structure which consists of support rings made of carbon fiber composite material located at the end of each cylinder, and which, in turn, are attached to a main support disk at each end of the tracker. Each of the 6 cylinders extends the full length of the tracking volume. The inner two superlayers of fiber detectors do not extend to the end of their respective cylinders but stop along an aperture line (at $|\eta| = 2.3$) defined by the next superlayer. This implies that a track, with rapidity $|\eta| \leq 2.3$, will traverse no less than three superlayers.




FIG. 4-59. Close-up view of the support structure for the outer tracker showing the interface between the outer tracker and the silicon system. The routing of waveguide fiber bundles from the tracking elements is shown. No bending radii are smaller than 10 cm.

The silicon tracking system support is also shown in Fig. 4-59. A removable cylindrical drum attaches to the main support disk. The attachments are such that the silicon tracking system can be inserted or removed independently of the outer tracking system. The large support disks are the backbone that will support the entire tracker. These have been designed to provide large access spaces for the bundling and routing of waveguide fibers which transmit the scintillation light to the VLPC's situated outside of the calorimeter.

The fiber ribbons extend from the center of the detector ($z = 0$) to the $|\eta| = 2.3$ line for the inner layers (B1-B3) and to the end of the tracking volume, $z = 4.3$ m, for the outer layers (B4-B6). Hence, for a given superlayer, the longitudinal coverage is split in the middle, with fibers read out at each end of the detector. The resultant crack at the center of the detector is of very small dimension (few mm). The ribbons of scintillating fiber are optically coupled/spliced to clear fiber waveguides which transmit the produced scintillation light to the VLPC's (Fig. 4-60). For the fiber waveguides, the path length varies from a minimum of ~ 4.5 m to a maximum of ~ 7 m, with the longest corresponding to the path from the innermost barrel superlayer (B1) to the VLPC's. However, these long waveguide paths are coupled to scintillating fibers of shorter length as Table 4-21 indicates.

The VLPC's convert the scintillation light to amplified signals for detector readout. These structures are Si impurity- band conduction devices designed originally for infrared imaging, but now optimized for visible light response. VLPC's exhibit excellent quantum efficiency at visible wavelengths ($QE \sim 0.85$), have pulse rise times of order 5 ns and can tolerate magnetic fields of several Tesla without degradation in performance. The devices must be operated at cryogenic temperatures (6-9 K), and helium cryostats are required.

The 64 VLPC canister cryostats required for the fiber tracking system are located around the perimeter of the barrel calorimeter. Thirty-two canisters are at each end, as close as possible to the services crack between the barrel and endcap calorimeters. Since the end cap calorimeters are rolled out to access the interior of the central detector, all the canisters are supported off the barrel. In the current engineering design, the VLPC cryostat is of the "pool-boiling" type. To implement this approach, all canisters are arranged in a Ferris-wheel fashion, with each canister positioned so that elements can be removed vertically.

A sketch of one of the cryostats is shown in Fig. 4-61. Large scale integration is achieved by packaging 512 VLPC channels into modules which are housed in a card-like cassette as shown in Fig. 4-62. Each of the cryostats supports 16 of these cassettes, for a total of 8192 channels. The VLPC modules are maintained at the appropriate operating temperature (6-9 K) by conductive heat flow to a copper heatsink and by the continuous flow of gaseous helium across each module. Fine temperature control is achieved by a small heater and two temperature monitors.

The portion of the cassette that extends above a mounting flange is at ambient temperature and houses the electronics and fiber-optic couplers for the cryostat. The VLPC outputs (cold end) are electrically connected to preamplifiers situated on PC boards, located in the warm end of the cassette, using low capacitance ribbon cables. The PC boards contain the amplifier/shaper/discriminator circuitry, a temperature monitor, and control circuitry.

One of the design criteria for the VLPC cryostats is to accurately control the VLPC temperature while minimizing the liquid helium consumption and optimizing helium gas recovery. The heat load due to conduction, radiation, VLPC heating, and Joule heating is estimated to be 0.35 mW per channel under normal operating conditions. This translates to a 175 W heat load for all 475 K channels, which should pose no problem.

A summary of the physical parameters of the VLPC Cassette/Cryostats is presented in Table 4-22.




FIG. 4-60. Isometric view of the tracker and calorimeter showing the placement of VLPC cryostats and sample fiber waveguides which "pipe" the scintillation light from the tracker to the VLPC's.

4.8.3. Material in the tracking volume

The material for the combined silicon and scintillating fiber tracking system, expressed in radiation lengths, is displayed as a function of η in Fig. 4-63. For the fiber portion of the tracker, the material is dominated by the scintillating fibers, with a ratio of fiber material to support structure of 3:1. Averaged over the interval $|\eta| < 2.3$, the full tracker represents 18.7% of a radiation length, discounting the outer fiber superlayer which is positioned right against the magnet. Since superlayers B1, B2, and the axial layers of B3 provide the trigger, the material outside these elements should be ignored for trigger purposes. This is illustrated in Fig. 4-63, in which the dashed curve indicates the material affecting the trigger, which, averaged over the interval $|\eta| < 2.3$, amounts to 12.3% of a radiation length.




FIG. 4-61. VLPC cryostat supporting 8,192 channels.




FIG. 4-62. VLPC cassette supporting 512 channels.

Table 4-22
VLPC cassette/cryostat: physical parameters.

A. Footprint	20" diameter by 20" height
B. Orientation	Vertical to allow pool boiling helium option
C. Subassembly size	Replaceable unit containing 512 pixels
D. Assembly size	64 Assembly units containing a maximum of 8,192 VLPC channels
E. Reliability	300 years mean time to failure per pixel
F. Useful lifetime and number of thermal cycles	10 year lifetime 50 thermal cycles per subassembly
G. Cool down time	6-8 days for cool down of a full assembly
Replacement time	3-4 hours cool down for swapping one subassembly (subassembly capable of being withdrawn cold)
H. Environmental factors	Room temperature helium recovery required.




FIG. 4-63. Material in radiation lengths for the combined silicon and fiber tracker as a function of η .

4.8.4. System performance

The motivation for building a scintillating fiber outer tracker is to provide a system capable of performing pattern recognition, track reconstruction, and Level 1 triggering in the crowded SSC environment, for luminosities at and well beyond the design value.

Track reconstruction

For the silicon/scintillating fiber tracker, the performance is comparable to that of the baseline tracking system at design luminosity (see Figure 4-6 and Table 4-5), and should exceed the performance of that system as the luminosity increases. This derives from the very low occupancy and extensive η coverage of the fiber superlayers, and the ability to resolve individual beam crossings with this detector. Table 4-23 summarizes the average occupancies in the fiber superlayers at design luminosity.

Table 4-23
Average occupancies at SSC design luminosity.

Superlayer	Radius (cm)	$\Delta\eta$ coverage	Occupancy
B1	60.0	2.3	0.017
B2	76.0	2.3	0.013
B3	92.0	2.3	0.012
B4	108.0	2.1	0.006
B5	136.0	1.9	0.003
B6	165.0	1.7	0.002

Pattern recognition

A substantial pattern recognition capability has been demonstrated for the silicon/scintillating fiber tracker for both isolated tracks and tracks within jets of p_t up to 1 TeV/c [45,46]. One can gain an intuitive feel for this capability by noting that there is a significant commonality in the two systems: silicon strips and scintillating fibers are topologically similar and the output from both systems is strictly digital (latch) information. Fibers are simply long flexible strips. Thus it is possible to apply common analysis techniques to both systems. For example, hits in adjacent pairs of superlayers in the barrel region can be combined with a vertex constraint using a simple circle-point fit to form vectors in curvature-phi space and the error ellipse for each vector can be plotted [47]. Curvature (K) is proportional to (p_t^{-1}) and ϕ_0 is the azimuthal angle of the emergent particle. Vectors from different superlayer-pairs associated with the same track will

lie on top of one another forming clusters in this space. Track candidates of high transverse momentum, such as leptons from Higgs decay, have near-zero curvature and are easily located. Because the error ellipses of the vectors in the silicon and fiber systems are of comparable size, high p_t leptons appear as well isolated clusters in (K, ϕ_0) space for both systems. This implies that one can make a uniform and effective application of the vector clustering technique for the combined silicon and fiber tracker.

Trigger concept

The SDC performance requirements call for a first level tracking trigger with $\Delta p_t/p_t^2 = 10 \text{ (TeV/c)}^{-1}$ with a factor of two improvement at Level 2. The efficiency is required to be 96% per track with fewer than 0.05 false triggers per calorimeter bin.

The scintillating fiber tracking trigger is formed from the three inner superlayers B1, B2, and B3. Each superlayer is made up of two double-layers of fibers spaced so that each doublet may be considered as an independent measurement for triggering, and the superlayers produce local vectors. We calculate the trigger performance assuming that, for trigger purposes, each doublet is read out in bins of one fiber width [6,48], and the superlayers are linked to form track candidates.

The momentum resolution of the fiber tracking trigger can be calculated from the radii and resolutions of the inner and outer trigger layers since the intermediate layer serves only to confirm the existence of a track. Assuming a radial vertex constraint the resolution is $\Delta p_t/p_t^2 = 6.0 \text{ (TeV/c)}^{-1}$. This value easily satisfies the Level 1 trigger requirement and is very close to the value required for a Level 2 trigger. At Level 2, the silicon information is expected to be available and can improve this measurement significantly. Figure 4-64 shows the trigger efficiency (turn-on curve) as a function of p_t for a 10 GeV/c threshold. In fact, the trigger will be reported using four bits (3 bits of slope and a sign bit), so that the threshold may be adjusted as required.

The usefulness of the tracking trigger at Level 1 depends on the rate of false triggers. These arise from two sources: real tracks with momenta below threshold which are accepted due to finite momentum resolution, and hits from different low- p_t tracks which accidentally line up to mimic a high- p_t track. The yield per beam crossing of false high p_t tracks is shown in Fig. 4-65 for various options [6]. For trigger logic in which two out of the three inner superlayers are required, the fake trigger rate reaches 0.4/crossing at design luminosity and 450/crossing at ten times design luminosity. Given the calorimetric granularity of 128 bins at Level 1, there is adequate fake rejection at design luminosity. At ten times design luminosity, good rejection is achieved by including additional superlayers in the trigger. This would imply a more restrictive η coverage, if layers B4, B5, or B6 are used.

Electronics implementation

Architecturally, the electronics for the Scintillating Fiber Tracker can be viewed as a multi-layered structure. It is shown schematically in Fig. 4-66.

- a) The first layer provides the interface and support for the VLPC's. This includes VLPC read-out, the supply of VLPC bias voltages and temperature monitoring, and power for the amplifier/shaper/discriminator circuits within the VLPC cassettes.
- b) The second layer provides trigger delay storage for both Level 1 and Level 2.
- c) The third layer provides pipelined track information which is reported to the SDC trigger system for linking with information from the calorimeter and muon detectors.
- d) The final layer provides the readout path to the DAQ system and also provides for low-rate diagnostic monitoring of the system.

Functionally, the electronics is implemented in 16 crates which are mounted on the calorimeter close to the VLPC cryostats. Each crate contains a clock/trigger card, a CPU, and up to 16 readout cards on which are implemented all four of the functions described above. The system as currently designed uses two ASIC's. The first of these, the segment-finder ASIC, provides all Level 1 and Level 2 pipelined data




FIG. 4-64. Trigger efficiency as a function of p_T for the scintillating fiber tracking trigger using superlayers B1, B2, and B3.




FIG. 4-65. Accidental trigger rate as a function of luminosity for the scintillating fiber tracker: using superlayers B1-B3 and requiring that 2 out of 3 participate in trigger; using superlayers B1-B4 and requiring that 3 out of 4 participate; using superlayers B1-B5 and requiring 4 out of 5 participate.

storage, provides segment finding within a superlayer, and the DAQ interface for 64 fibers. The second segment-linking ASIC correlates the found segments as determined by the first ASIC to generate a trigger based upon the stiffness of a track candidate.

The algorithm for track segment finding uses a fixed pattern of combinatorial "AND/OR" logic to find high p_T hits in a given superlayer with a resolution of one fiber diameter. The segment-linking ASIC serves to correlate track segments from three to six superlayers within a narrow angular cone corresponding to that expected for a track of 10 GeV/c or greater transverse momentum.

4.8.5. Development of fiber tracking technology

To realize the potential of scintillating fiber tracking technology, an ongoing program of R&D has resulted in significant technical developments in the following areas: high efficiency, low-loss, radiation-resistant fiber scintillators and waveguides; high quantum efficiency, visible-light-sensitive photon counters (VLPC's) and associated front-end preamplifiers; and demonstration of fiber tracking in bench and beam tests.

Scintillator and waveguides

For demonstration of proof-of-principle for fiber tracking and for selection of suitable candidate materials for construction of the SDC fiber tracker, we have chosen scintillating and nonscintillating fiber waveguides having polystyrene (PS) cores and polymethylmethacrylate (PMMA) claddings. This choice has been motivated by several factors: (1) PS scintillators have been produced which have good fluorescence efficiency and fast decay; (2) clear fiber waveguides based on these materials exhibit excellent optical attenuation length over long distances for visible wavelengths; (3) PS scintillators have shown good resistance to radiation damage; and (4) there is substantial experience in working with these materials. Fibers and waveguides have been drawn with excellent dimensional tolerance ($\leq 1\%$ on the diameter) by several commercial houses.

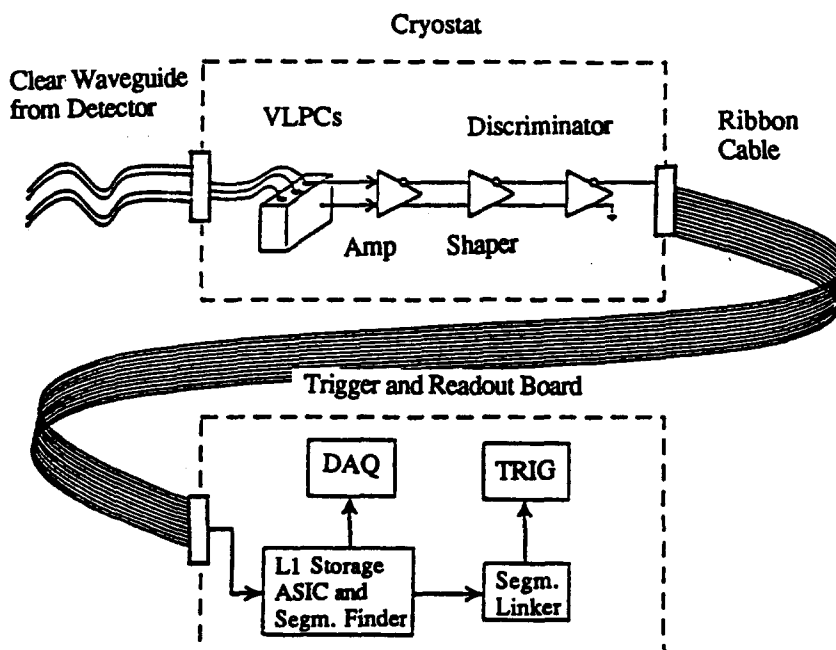


FIG. 4-66. Schematic of the scintillating fiber tracking front end electronics.

Substantial effort has been devoted to the development and identification of highly efficient scintillation compositions in polystyrene. These studies indicate that a scintillator which incorporates a secondary dye material exhibiting intramolecular proton transfer (IPT) is the preferred choice. Several organic fluor systems exhibiting this IPT property have been studied, including the hydroxybenzothiazoles (HBT) and 3-hydroxyflavones (3HF). Of these, the dye which has provided the best performance to date has been 3HF, incorporated into polystyrene with the primary dye p-terphenyl (PTP). This material has been established as the "standard" scintillator for the fiber tracker and boules of this composition have been drawn by several commercial vendors into fibers of varying diameters for studies.

The standard material has fluorescence emission peaked near 530 nm, reasonable speed (5–7 ns), and acceptable efficiency [50]. It exhibits little optical self-absorption because the secondary (3HF) is an IPT fluor and because this dye is incorporated into the polystyrene at low concentration ($\leq 10^3$ ppm). The long-wavelength emission ($\lambda = 530$ nm) is well matched to two important characteristics of polystyrene: general optical transmission is best in unirradiated polystyrene for visible wavelengths $500 \text{ nm} < \lambda < 700 \text{ nm}$ (Fig. 4-67) [6,51] and the effects of radiation damage on optical transmission in polystyrene are substantially reduced for wavelengths $\lambda > 500 \text{ nm}$ [52].

Scintillating fibers fabricated of standard material in the core and acrylic (PMMA) for the cladding exhibit long optical attenuation length ($\sim 4 \text{ m}$). Additionally, fibers of the same base composition, but without the scintillation dyes, possess attenuation lengths of $\sim 9 \text{ m}$ in the wavelength range of interest [6]. Such waveguides are to be used to transmit the produced scintillation light to photodetectors situated outside the tracking volume.

The scintillation and waveguide fiber materials exhibit good resistance to radiation damage. We have irradiated several samples of the base material in the form of small boules of 1 cm diameter and 10 cm length. Figure 4-68 shows the results of a 1 Mrad dose from ^{60}Co , corresponding to 50 years of exposure at SSC design luminosity at the inner tracking radius of $r = 60 \text{ cm}$. As can be seen in the figure, the sample shows negligible transmission loss, after annealing, for wavelengths above $\lambda \geq 500 \text{ nm}$. Irradiation was performed in a nitrogen atmosphere, followed by annealing in oxygen. It is only at elevated levels of irradiation that one begins to observe transmission loss in the base material, after annealing, for wavelengths above $\lambda \geq 500 \text{ nm}$. At a 10 Mrad dose (500 SSC year equivalent), one sees a $\sim 5\%$ loss in transmission. At




FIG. 4-67. Absorption in unirradiated polystyrene as a function of wavelength (solid curve) and fluorescence emission of the standard scintillation material, PS/PTP/3HF (dashed curve).




FIG. 4-68. Transmission in standard PS/PTP/3HF scintillator as a function of wavelength, before and after irradiation by ^{60}Co to 1 Mrad, and after annealing.

these very high doses, a modest reduction ($\sim 20\%$) in scintillation efficiency is also observed [53].

Significant activity is now in progress to measure the performance of scintillating and waveguide fibers irradiated at Fermilab. In this program, fibers and fiber ribbons of up to 4 m length are subjected to realistic doses and dose rates of ionizing radiation under different ambient conditions. Optical properties of the irradiated samples are now being compared with unirradiated samples maintained in similar gas atmospheres, but well away from ionizing radiation. Based upon these comparisons, fiber/ribbon degradation and annealing are being studied as a function of exposure, time, and atmosphere. Preliminary measurement of the optical attenuation length as a function of exposure rate and atmosphere for the standard composition indicates a $< 5\%$ drop in transmission over 4 m of scintillator in an exposure equivalent to 3 standard SSC years at the inner tracking radius of 60 cm, followed by annealing in air.

Photodetectors

Visible light photon counters (VLPC's) combine good quantum efficiency with high gain and low gain dispersion to achieve single-photon counting capability, which is essential to record the signals from minimum ionizing particles. The development and refinement of these devices is in progress at Rockwell International Science Center and Electro-Optical Center.

Several lots of VLPC devices have been successfully produced and Fig. 4-69 displays measured or expected quantum efficiency for these structures as a function of incident wavelength. It should be noted that the visible-light quantum efficiency of VLPC's is an order of magnitude better than that of a conventional photomultiplier at visible wavelengths. The first lot designated HISTE-I has a measured quantum efficiency of $\sim 85\%$ at a wavelength $\lambda = 565$ nm, has been produced in linear arrays, and has been delivered in significant quantity (~ 100 channels). However, these devices have high infrared (IR) quantum efficiency, and because of military restrictions (ITAR), have been available only for bench tests at universities. In the second lot (HISTE-II), the infrared response has been significantly reduced, with a measured QE $\sim 0.05\%$ for wavelengths $\lambda \geq 1$ μm . Use of these devices is unrestricted, and 32 channels of these detectors were successfully operated in beam tests in the New Muon Laboratory at Fermilab during the Fixed Target Run, Dec. 91-Jan.92. The third lot (HISTE-III) has been produced recently, and exhibits high sensitivity to visible light (QE $> 75\%$) and low infrared sensitivity (QE $< 1\%$). These devices are to be delivered in substantial quantity for performance and characterization studies.




FIG. 4-69. Spectral quantum efficiency for VLPC devices. HISTE-I are DoD restricted devices. HISTE-II and III are unrestricted devices.

In the SDC configuration, the VLPC's will be situated in cryostats outside the calorimeter where radiation levels are low. Based on Rockwell's experience with radiation effects in similar devices (SSPM's), it is estimated that the thresholds at which noticeable performance degradation occurs for VLPC's are: $> 10^5$ rad for gamma radiation and $> 10^{11}$ n/cm² for neutrons. Table 4-24 summarizes the geometric and operating characteristics of VLPC's [54].

Table 4-24
Geometric and operating characteristics of visible light photon counters.

Parameter	Achieved	Goal
Active area	$875 \times 875 \mu\text{m}^2$	1 mm diameter
Array configuration	1×8	1×10
Contact resistance	1000 Ω	$< 100 \Omega$
Pulse rise time	< 5 ns	< 5 ns
Average gain	20,000	30,000
Gain dispersion	$< 30\%$	$< 30\%$
Effective QE at 565 nm	85%	70-90%
Infrared QE at 15	$< 0.5\%$	$< 2\%$
Dead time	None (continuous)	None (continuous)
After pulses	$< 0.01\%$	$< 0.01\%$
Dark pulse rate	10 kHz	5 kHz
Saturation pulse rate	25 MHz	25 MHz
Average current	< 200 nA	< 200 nA
Average power	1.4 $\mu\text{W}/\text{channel}$	1.6 $\mu\text{W}/\text{channel}$
Breakdown voltage	7.5 V	> 8.5 V
Operating bias voltage	6-8 V	6-8 V
Neutron damage level	$> 10^{11}$ n/cm ²	$> 10^{11}$ n/cm ²
Operating temperature	6-8 K	6-8 K

Preamplifiers

Recent studies have shown that several types of fast, low-noise preamplifiers can be used successfully to read out VLPC's [6,55]. These include Fermilab chips QPA02 and VTX, developed for readout of silicon microstrip detectors and CDF vertex chambers, respectively; and the Penn TVC chip, developed initially for readout of straw tubes and muon chambers for SDC. Pulse rise of < 5 ns has been observed and excellent separation of single photoelectron signals from noise has been demonstrated.

Tests of fibers and VLPC's

Several fiber tracker tests have been performed: bench tests of fibers with HISTE-I VLPC's using cosmic rays, and beam tests of fibers with HISTE-II VLPC's and a multianode photomultiplier in the muon beam at Fermilab. In the bench tests, the performance of arrays of 16 scintillating fiber elements has been studied. The scintillating fibers were configured in a lattice of 4 ribbons containing 4 fibers each. Kuraray fibers of standard composition, $830\ \mu\text{m}$ in diameter, were close-packed in center-to-center spacing to form the ribbons. Each scintillating fiber was optically coupled to a clear waveguide fiber of the same diameter, and the fiber/waveguide elements were read out individually by HISTE-I photodetectors. The signals were amplified by transimpedance amplifiers operated at room temperature. Figure 4-70 displays a typical pulse height spectrum obtained from a fiber excited by a ^{60}Co source. Individual lines are clearly resolved for as many as 7 simultaneously detected photoelectrons.

Figure 4-71 shows the mean detected photoelectron yield for various positions of the cosmic ray trigger telescope. At the position of the trigger counters farthest removed from the photodetector (corresponding to 3.5 m of scintillating fiber and 3 m of clear waveguide), we observe a mean detected yield of 6.5 photoelectrons with the ends of the scintillating fibers mirrored, and ~ 5 photoelectrons with the ends unmirrored.

We propose to implement the SDC fiber tracker with unmirrored scintillating fibers of $925\ \mu\text{m}$ diameter coupled to waveguide fibers of 1 mm diameter. Unmirrored fibers have half the occupancy of mirrored fibers, due to the finite propagation time for light in the material (5 ns per meter length of fiber). Although direct measurements of the photoelectron yield from fibers of these diameters have not yet been performed, it is straightforward to estimate the expected detected light yield for the proposed SDC geometry, scaling from the measured values for $830\ \mu\text{m}$ fibers. The mean photoelectron yield for $925\ \mu\text{m}$ diameter fibers at $\eta = 0$ is expected to be ~ 4.5 , which should assure efficient doublet ribbons of fibers and allow us to achieve our performance goals for triggering and tracking.

Beam tests have been performed at Fermilab (T839) in the New Muon Laboratory using ribbons of fibers read out with HISTE-II VLPC's and a Hamamatsu multianode photomultiplier tube. Objectives of the test included the recording of high energy tracks in a system of fiber superlayers, the implementation of a new small cryostat system capable of supporting 32 VLPC channels, and the implementation of fast, room-temperature preamplifiers in conjunction with the VLPC's.

Particle trajectories were recorded in four superlayers of fibers. The first of these was a 32 element array of Kuraray fibers of $830\ \mu\text{m}$ diameter and standard composition arranged as four 8-fiber layers, staggered at $1/4$ fiber diameter spacing, and readout with HISTE-II VLPC's. The remaining layers consisted of a 96 element array of Bicron BCF-10 (blue) fibers arranged in three 32 element superlayers with $1/4$ diameter stagger, and readout with the multianode photomultiplier tube.

For this test, a compact cryogenic system consisting of a dewar and VLPC cassette cryostat of Rockwell design was prepared. The miniaturization of the system allowed for reduced capacitance and the effective implementation of room temperature, fast preamplifiers. In these studies, VTX chips were used. As Fig. 4-72 shows, the signals from this HISTE-II/VTX system exhibited excellent gain, signal-to-noise, and rise-time (< 5 ns) performance. This result demonstrates that room temperature preamplifiers will operate effectively with VLPC's and obviates the need for low temperature preamplifiers inside the dewar.

Numerous tracks of high energy muons were recorded during the test run. The measured track resolution derived from these studies indicates that the spatial resolution for a double-layer of fibers is




FIG. 4-70. Typical pulse height spectrum from a standard fiber readout with a HISTE-I VLPC, exhibiting photopeaks and excellent gain dispersion.




FIG. 4-71. Average number of detected photoelectrons per fiber layer in the cosmic ray tests using HISTE-I VLPC photodetectors. The values on the abscissa indicate the physical separation in meters of the trigger counter telescope from the VLPC's, and includes 3 m of clear fiber-optic waveguide, an optical splice, and then up to 4 m of scintillating fiber. Squares and X's are for mirrored end; diamond and cross are for unmirrored end.




FIG. 4-72. Oscilloscope trace showing the signals from the Fermilab VTX warm amplifier connected via low capacitance ribbon cables to a HISTE-II VLPC operated at 9 K. The vertical scale is 5 mV/div, the horizontal scale is 10 ns/div. Notice that the single photopeak is clearly visible above the noise. Rise time is of the order of 5 ns.

$\leq 120 \mu\text{m}$, consistent with our expectations [6].

The cosmic ray and beam tests have demonstrated the potential of the fiber tracking technique: that VLPC's and room temperature preamplifiers can meet the response requirements of the SDC, that a compact cryostat system can be quickly cooled and reliably operated, and that system operation is remarkably straightforward and stable.

4.8.6. R&D Schedule

Table 4-25 itemizes the R&D objectives for FY92 and FY93. Among the principal elements of this program are system tests of 1 K and 10 K fiber and VLPC channels.

Table 4-25
R&D milestones for the scintillating fiber tracker FY92-FY93.

Scintillator R&D	
Selection of "Standard" composition	Jan. 92
Characterization of other candidates	Sep. 92
Waveguide R&D	
Selection of waveguide composition	Jan. 92
Selection of cladding material	Sep. 92
Splicing R&D	
Review of splicing techniques	May 92
Prototype candidate techniques	Jul. 92
Selection of splicing technique and sequence	Sep. 92
Fiber manufacturing engineering	
Optimization of fabrication process	Sep. 92
Development of automated Q/A apparatus	Jan. 93
Production of significant quantities of fiber	Feb. 93
Ribbon Manufacturing Engineering	
Selection of ribbon fabrication technique	Sep. 92
Engineering design of production facility	Jan. 93
Determination of Q/A procedures and equipment	Jan. 93
Readout R&D (VLPC)	
Test of 128 channel system	Jun. 92
Further IR desensitization	Sep. 92
Reduction in ohmic contact resistance	Sep. 92
Latching/preamplifier issues	Sep. 92
Warm preamp development	Sep. 92
Readout engineering (VLPC)	
Conceptual design for the readout system	Mar. 92
Engineering design for the readout system	Sep. 92
Operation of a 1 K fiber/VLPC system	Jan. 93
Fully functional pre-production prototype	Jun. 93
Operation of a 10 K fiber/VLPC system	Jan. 94
Ribbon placement onto Superlayers	
Stable-base cylinder prototype development and construction	
Thin cylinder	Jul. 92
Thick cylinder	Sep. 92
Engineering placement concept	Sep. 92
Construction of placement jigs and fixtures	Jun. 93
Placement of fiber superlayers onto prototype cylinder	Sep. 93
Trigger/ASIC Development	
Discrete technology prototype	Jun. 92

References:

1. W.T. Ford *et al.*, "Requirements for the SDC Tracking System", Solenoidal Detector Collaboration Note SDC-91-00127 (December 1991).
2. J.A.J. Matthews, "High P_t Forward Rapidity Tracking Trigger Using Silicon Planes", Solenoidal Detector Collaboration Note SDC-90-00067 (1990); H.-J. Trost, "Some Physics Requirements for Triggering in the Intermediate Tracking System of SDC", ANL/HEP/TR-91 (1990).
3. A. Weinstein *et al.*, "Silicon Tracking Conceptual Design Report", SCIPP 92/04 (January 1992).
4. B. Adrian *et al.*, "SDC Modular Straw Outer Tracking System Conceptual Design Report", Solenoidal Detector Collaboration Note SDC-91-00125 (January 1992).
5. B. Foster *et al.*, "Intermediate Angle Track Detector Conceptual Design Report", Solenoidal Detector Collaboration Note SDC-92-171 (January 1992).
6. "Conceptual Design: Scintillating Fiber Outer Tracker," Solenoidal Detector Collaboration Note SDC-92-174 (Jan. 1992).
7. B. Hubbard, "B Jet Tagging with the SDC Tracking System," SDC Note 92-212.
8. D. Adams *et al.*, "Report of the Task Force on Impact of Material in Tracking Volume," Solenoidal Detector Collaboration Note SDC-91-00120 (1991).
9. Y. Unno, "The effect of material in the tracking volume - E/p and position matching," Solenoidal Detector Collaboration Note SDC-91-136 (December 1991). B. Hubbard and W. Lockman, "Effect of Material on Electron Tracking," SDC Note 91-109.
10. G. Sullivan and M. Miller, "Calorimetry Trigger Rates at the SDC," Solenoidal Detector Collaboration Note SDC-91-00099 (1991).
11. T. Gorski, J. Lackey, and W. Smith, "SDC Trigger Preliminary Conceptual Design," Solenoidal Detector Collaboration Note SDC-91-00089 (1991).
12. T. Gorski, J. Lackey, and W. Smith, "The First Four Microseconds: SDC Level 1 Pipeline Length," Solenoidal Detector Collaboration Note SDC-91-00139 (1991).
13. T. Gorski, J. Lackey, and W. Smith, "SDC Global Level 1 Processor: Clock and Control," Solenoidal Detector Collaboration Note SDC-91-00090 (1991).
14. J. Chapman, A. Dunn, and J. Mann, "SDC Straw Trigger Preliminary Conceptual Design," UM-HE-91-39 (1991).
15. J. Chapman and J. Mann, "Triggering in the SDC Using Mean Timer Synchronizers with Straw Drift Tubes," *Conference on Electronics for Future Colliders*, LeCroy Corporation (1991).
16. A. Seiden, "Systematic Error and Alignment for Barrel Detectors," SCIPP Preprint 91/20 (1991).
17. D. Groom, editor, Radiation Levels in the SSC Interaction Regions, "SSC Central Design Group Task Force Report", SSC-SR-1033 (June 10, 1988).
18. H. F.-W. Sadrozinski, SCIPP 91/01.
19. R. Openshaw *et al.*, "Etching of Anode Wire Deposits with CF_4 /Isobutane (80:20) Avalanches", *Proceedings of the Symposium on Detector Research and Development for the Superconducting Super Collider*, Fort Worth, Texas, Oct. 15-18, 1990, (World Scientific, Singapore, 1991) p. 231.
20. J. A. Kadyk, J. Va'vra and J. Wise, *Nucl. Inst. Methods A300* (1991) 511.
21. Rick van Berg *et al.*, Workshop on Major SSC Detectors, Tucson, AZ, February 18-23, 1990.
22. H. Ikeda and N. Ujiie, *Nucl. Instrum. Methods A281* (1989) 508.
23. E.J. Kennedy *et al.*, "Radiation Effects on JFETS, MOSFETS, and Bipolar Transistors, as Related to SSC Circuit Design," *Proceedings of the Symposium on Detector Research and Development for the Superconducting Super Collider*, Fort Worth, Texas, Oct. 15-18, 1990, (World Scientific, Singapore, 1991.)

24. N. Cartiglia *et al.*, SCIPP Preprint 92/05 (1992).
25. W. Dabrowski *et al.*, SCIPP Preprint 91/24 (1991); S. Tedja *et al.*, University of Pennsylvania preprint, submitted to Nucl. Instrum. Methods.
26. H. Ziock, *IEEE Trans. Nucl. Sci.* NS37 (1990) 1241.
27. J.F. Arens, "Designs for a High Energy Physics Detector Array", UCB/SSL internal report, 1989.
28. D.R. Nygren, "Silicon Tracking Devices for the SSC", *Proceedings of the Summer Study on High Energy Physics in the 1990's*, June 27-July 15, 1988, Snowmass, Colorado, edited by S. Jensen (World Scientific, Singapore, 1989).
29. O. Barkan and G. Atlas, "Conceptual Design Review", Hughes Aircraft Co., May 2, 1990.
30. D.R. Nygren, "Concept for Pixel Array Readout with Good Time and Charge Resolution", LBL internal report, August 21, 1991.
31. D.R. Nygren, "Converting Vice to Virtue: Can Time-Walk be Used as a Measure of Deposited Charge in Silicon Detectors?", LBL internal report, May 22, 1991.
32. D.G. Cassel, G.G. Hanson, *et al.*, *Proceedings of the 1986 Summer Study on the Physics of the Superconducting Supercollider*, edited by R. Donaldson and J. Marx, Snowmass, CO, 1986, p. 377.
33. J. Florent, J. Gaudaen, L. Ropelewski, and F. Sauli—CERN (to be published).
34. A. Oed Nucl. Instr. Methods A263, 351 (1988).
35. A. Oed *et al.* Nucl. Instr. Methods A284, 223 (1989).
36. F. Angelini *et al.*, Nucl. Instr. Methods A283, 755 (1989).
37. F. Angelini *et al.*, "The Microstrip Gas Chamber," INFN PI/AE 90/6 (1990).
38. F. Angellini *et al.*, "Results from the first use of Microstrip Gas Chambers in a High-Energy Physics Experiment," CERN-PPE/91-122 (1991).
39. F. Angellini *et al.*, "A Microstrip Gas Chamber on a Silicon Substrate" INFN PI/AE 91/10 (1991).
40. F. Hartjes *et al.*, "A Prototype Microstrip Gas Detector," Nucl. Instr. Methods A289, 384 (1990).
41. Geilberts *et al.*, "Tests of the Performance of different gas mixtures in microstrip gas counters," NIKHEF-H/ 91-23 (1991).
42. E.F. Barasch *et al.*, "Recent Developments in Gas Microstrip Chambers," *Proc. 5th Pisa Meeting on Advanced Detectors* (1991).
43. R. Bouclier *et al.*, "Microstrip Gas Chambers on Thin Plastic Supports," IEEE, Nuclear Science Symposium, Santa Fe, November 2-9, 1991.
44. M. D. Petroff and M. G. Stapelbrock, *IEEE Trans. on Nucl. Sci.* NS36, 158 (1989); and M. D. Petroff and M. Atac, *IEEE Trans. on Nucl. Sci.* NS36, 163 (1989).
45. D.L. Adams, G. Eppley and J. Skeens, "Pattern Recognition in a Silicon and Scintillating Fiber SDC Tracking System," Solenoidal Detector Collaboration Note SDC-91-131 (November 1991).
46. D.L. Adams, "Pattern Recognition in a Silicon and Scintillating Fiber Tracking System II," SDC-note, in preparation.
47. J. Hylen, private communication (December 1991).
48. D.L. Adams, "Momentum Smearing and False Trigger Rates in a Scintillating Fiber Tracking Trigger," Solenoidal Detector Collaboration Note SDC-91-130 (November 1991).
49. R. Lewis *et al.*, "Evaluation of a Proposed CDR Version of the Fiber Tracker, private communication (February 1991).
50. B. Abbott *et al.*, "Scintillating Fiber Detectors," *Proceedings of the Symposium on Detector Research and Development for the Superconducting Super Collider*, Fort Worth, TX (1990), pp. 90-99.

51. J. Allitti *et al.*, *IEEE Trans. Nucl. Sci.* **36** (1989), pp. 29-34.
52. A. Bross and A. Pla-Dalmau, American Chemical Society Symposium Series, October 1990 and Fermilab/PUB 90-224.
53. A. D. Bross *et al.*, *IEEE Trans. Nucl. Sci.* **NS38**, 994 (1991). See also Conceptual Design Report: Scintillating Fiber Outer Tracker, Draft Version (Nov. 12, 1991)—SCT 000004.
54. VLPC Cassette/Cryostat Conceptual Design Review, Rockwell International Science Center, Anaheim, CA (February 1992), in press.
55. A.E. Baumbaugh *et al.*, Submitted to the Proceedings of the 1991 IEEE Nuclear Science Symposium, Santa Fe, NM (November 1991).

5. Superconducting magnet

5.1. Introduction

The SDC solenoid is designed to provide an axial magnetic field of about 2 T over the tracking volume. Particle detectors within this volume will measure the trajectories of charged tracks emerging from the 40 TeV center-of-mass collisions. Particle momenta are determined from the measured curvature of these tracks in the field. The coil is surrounded by electromagnetic and hadronic calorimeters used to measure the energy and impact points of the particles. The magnetic flux return of the magnet is provided by the iron absorber plates of the hadron calorimeters. Since all wide angle particles must pass through the superconducting coil and cryostat before impacting the calorimeter, the desired calorimeter performance requires the material in the coil to be minimized both in terms of radiation lengths (X_0) and absorption lengths (λ_I). In practice, this requirement must be balanced against the overall requirement that the superconducting coil must be robust, predictable, and have excellent long term operational reliability. The SDC solenoid achieves a transparency of $1.2 X_0$ or $0.25 \lambda_I$ at $\eta = 0$. We believe that technical advances in thin aluminum stabilized superconducting magnets make such a magnet feasible and describe the design of such a magnet below. General requirements for the SDC solenoid are summarized in Table 5-1.

Table 5-1
General requirements of SDC solenoid.

Magnet envelope		
Cryostat	Inner radius	1.70 m
	Outer radius	2.05 m
	Total half length	4.389 m
Nominal magnetic field		2 T
Transparency	$(\eta=0)$	$1.2 X_0$
		$0.25 \lambda_I$
Cool down time		< 14 days
Quench recovery time		< 4 hr

5.2. Magnetic field design

Field uniformity and coil forces are strongly influenced by the proximity of the iron/scintillator HAC1 and HAC2 calorimeter elements. These effects were calculated in a series of axisymmetric finite element analyses using the ANSYS finite element program.

A 2 T solenoid with a 4 m tracking half length was modeled. Calorimetry segmentation and slot geometries were based on the calorimeter conceptual design report [1]. Each slot and iron plate was explicitly modeled with finite elements to avoid the approximation of "smearing" the B - H properties to account for the small volume of air in the calorimetry. Axial and radial decentering forces were found by perturbing the axial coil position (axial decentering) or the radial iron position (radial decentering) and comparing the resulting coil forces with those on a coil with coincident geometric and magnetic centers.

For most of these calculations, the endwall of the calorimeter was assumed to be located at 47 cm relative to the end of the current sheet. The actual location may be in a region of 57–64 cm. Fig. 5-1 shows the coil compressive force and the axial decentering force as a function of the endwall position. Fig. 5-2 shows axial field contours for the 47 cm endwall locations. Figure 5-3 shows $\int B \times dl$ as a function of pseudorapidity.

contributions to the transparency.

The safety, reliability and predictability are important, as is the thinness of the magnet. The following guidelines have been incorporated into the baseline coil design:

- To homogenize the stresses in the coil outer-support-cylinder cold mass, the mechanical design will limit the strain in the outer support cylinder to 0.1%.
- To eliminate unacceptable thermal stress in the cold mass following a quench, the maximum hot spot temperature will be limited to 100 K.

A useful parameter to gauge the relative transparency of superconducting coils is the ratio of stored magnetic energy to cold mass (E/M) [3]. The E/M ratio in the coil determines the average coil temperature rise after a quench. If we assume uniform dissipation of all of the coil's stored magnetic energy in the coil after a quench, the E/M ratio should be approximately equal to the coil enthalpy at the coil temperature after the quench,

$$E/M = H(T_2) - H(T_1) \simeq H(T_2),$$

where H is the coil enthalpy, T_2 is the average coil temperature after the quench, and T_1 is the coil initial temperature. In solenoids presently in use at various high energy physics laboratories, the E/M ratio is around 5 kJ/kg, as shown in Fig. 5-4. This corresponds to an average coil temperature rise of 65 K after a quench. If the quench propagation speed is very fast and results in a homogeneous energy dump into the coil, we may still allow another 15 K to reach a coil average temperature of 80 K, at which the thermal stress is still negligible. That means we may allow an E/M ratio of 8 kJ/kg, corresponding to $T_2 = 80$ K. Including a contingency, a design goal of 7.5 kJ/kg has been decided upon for E/M in the SDC solenoid. That enables us to realize a thickness of $1.2 X_0$.




FIG. 5-4. Ratio of stored energy to cold mass for existing thin detector solenoids.

To achieve this thickness, an improvement in the mechanical strength of the conductor is required. Having the conductor at uniform high stress is a very efficient and reliable way to support the electromagnetic forces within the strain limit mentioned earlier. In our design, the goal for the yield strength of the pure aluminum stabilizer has been set as 67 MPa. This is based on recent R&D [3,4].

An effort to reduce the effective thickness of the vacuum vessel has also been made. Either honeycomb or isogrid for the outer vacuum wall would be very effective for this purpose. Based on R&D on isogrid and brazed honeycomb structures, either material could save about $0.2 X_0$ relative to solid aluminum. Taking these efforts into account, we arrive at the radiation lengths as a function of η shown in Fig. 5-5. The bump around at $\eta = 1.5$ is induced by the coil support.

Table 5-2
Baseline design parameters of SDC solenoid.

Dimensions:		
Cryostat	Inner radius	1.70 m
	Outer radius	2.05 m
	Half length	4.389 m
Coil	Effective radius	1.84 m
	Half length	4.12-4.18 m
Conductor	Thickness	44 mm
Outer support cylinder	Thickness	31 mm
Electrical parameters		
Central field		2.0 T
Nominal current		8,000 A
Inductance		4.6 H
Stored energy		146 MJ
Stored energy / cold mass		7.4 kJ/kg
Typical charging time		1 hour
Mechanical parameters		
Effective cold mass		20 tonnes
Total weight		25 tonnes
Radial magnetic pressure		1.6 MN/m ²
Axial compressive force		11 MN

Table 5-3
Transparency of the solenoid.

Element	Thickness [mm]	X_0	λ_0
Outer vac. wall	(Isogrid)	11	0.138
	(Honeycomb)	(7.1)	(0.080)
Outer rad. shield	2.0	0.022	0.0051
Outer sup. cylinder	31.0	0.348	0.0787
Superconductor			
(Al stab)	39.0	0.438	0.0990
(Nb.Ti/Cu)	2.9	0.181	0.0167
(GFRP)	3.1	0.016	0.0058
(Al Strip)	2.0	0.022	0.0051
Inner rad. shield	2.0	0.022	0.0051
Inner vac. wall	6.0	0.067	0.0152
Super-insulation	2.0	0.007	0.0023
Total	(w/ isogrid)	1.261	0.2651
	(w/ honeycomb)	(1.203)	(0.2510)

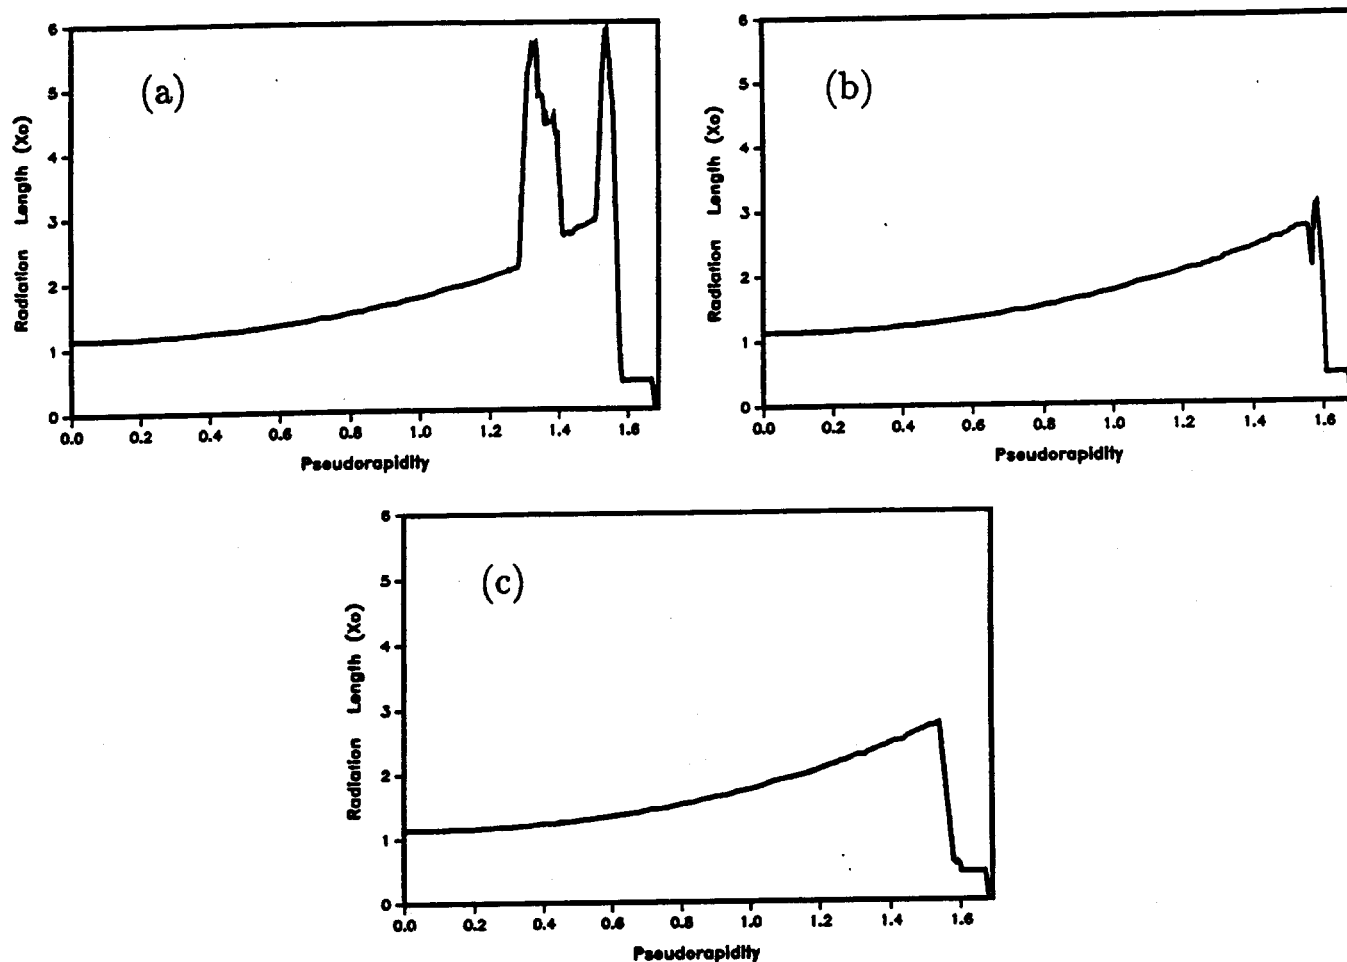


FIG. 5-5. Thickness of solenoid, in radiation lengths, as a function of pseudorapidity. (a) is a path through an axial support and is typical of 6% of the circumference, (b) is through the metallic end fitting of a radial support, 15% of the circumference, and (c) is through the GFRP section of a radial support, 21% of the circumference.




FIG. 5-6. Cross section of the end portion of the detector solenoid.

5.3.2. Coil design

The solenoid coil consists of a single layer-coil wound with aluminum-stabilized superconductor and an outer support cylinder as shown in Fig. 5-6. The coil is directly wound inside an outer support cylinder made of 5083-H32 aluminum alloy. Using an E/M ratio of 7.5 kJ/kg, a total cold mass thickness of 75 mm has been determined:

$$\begin{aligned}\frac{E}{M} &= \frac{E}{2\pi R t l \rho} \\ t &= \frac{E}{2\pi R l \rho (E/M)} \\ &= \frac{1.46 \times 10^5}{2\pi \times 1.85 \times 8.24 \times 2.7 \times 10^3 \times 7.5} \\ &= 0.075 \text{ m}\end{aligned}$$

Here R , t , l and ρ are the coil radius, total thickness including the support cylinder, coil length and the density of aluminum, respectively. The sharing of this total radial thickness between the superconductor and the support cylinder has been determined with two criteria: quench stability and safety, and conductor fabrication.

Even if the total stored magnetic energy were dumped entirely into the conductor, the conductor thickness of 44 mm (or equivalently, $E/M = 12.5$ kJ/kg) would still be adequate to keep the coil temperature at less than 100 K after a quench; this will be discussed later. With an operating current of 8000 A, the primary issue regarding conductor fabrication is the cross sectional aspect ratio of the conductor. A 44-mm thickness would result in a conductor width of 4.4 mm, assuming 2×0.1 mm turn-to-turn insulation. A conductor of this aspect ratio has been manufactured successfully and wound into a test coil as part of the R&D program. The outer support cylinder is therefore 31 mm thick, and as a result, a ratio of the conductor thickness to that of the support cylinder is about 3/2. The electrical ground insulation between the coil and support cylinder consists of two glass epoxy layers applied before the coil is wound. An electrical breakdown voltage of more than 2 kV is expected with an insulation thickness of 0.2 mm.

A helium cooling pipe with an inner diameter of 25 mm is welded to the outer surface of the support cylinder. The single serpentine path has 24 axial passes with a circumferential pitch of 45 cm, as shown in Fig. 5-7 and a total cooling path length of about 250 m. The cooling pipe is designed in accordance with the ASME piping code for the maximum pressure expected during the magnet's operation.

The solenoid is required to be very reliable during an initial test in air, and for long term operation at a nominal field of 2 T in the detector. Two mechanical stresses are especially important:

1. The combined hoop and axial stress, or stress intensity, in the superconductor during magnet operation in air. This generates a maximum shear stress ($2\tau = s_1 - s_2$) in the pure aluminum, at the axial coil center,
2. The shear stress at the boundary between the coil and the outer support cylinder due to the axial electromagnetic force, which is a maximum at the coil end.

First, we approximate the electromagnetic stress in the coil by using a simple analytic model. The radial magnetic pressure in a long solenoid is approximately

$$p_r = \frac{B_z^2}{2\mu_0} = 1.6 \text{ MPa} (= 227 \text{ psi}) .$$

Therefore, the hoop stress, s_ϕ , loaded in the coil and support cylinder is

$$\begin{aligned}s_\phi &= \frac{R}{t} \times p_r \\ &= 39 \text{ MPa}\end{aligned}$$




FIG. 5-7. Solenoid cold mass. Radial supports are on both ends, axial supports are on one end only.

where R is the coil radius and t is the total coil outer support cylinder thickness. The axial compressive force in a long solenoid in air is approximately

$$\begin{aligned} F_z &= dE/dz \\ &\simeq \frac{d}{dz} \left(\frac{B_z^2}{2\mu_0} S z \right) \\ &= 1,700 \text{ tonf} , \end{aligned}$$

where E is the stored energy, z is the effective coil length, and S is the area of the effective magnetic aperture. The axial magnetic force in the present design is reduced to $F'_z = 1100$ tonnes because the ferromagnetic calorimeter absorber is located close to the coil. Therefore, the practical axial stress in the coil and support cylinder is approximately

$$\begin{aligned} s_z &= \frac{F'_z}{2\pi R t} \\ &= -13 \text{ MPa} , \end{aligned}$$

where the negative sign indicates a compressive stress. These two stresses will be combined and accumulated at the coil center and the maximum combined stress in the conductor is given by

$$\begin{aligned} 2\tau &= (s_1 - s_2) + (s_2 - s_3) + (s_3 - s_1) \\ &\simeq (s_\phi - s_z) \\ &= 52 \text{ MPa} \quad (< \sigma_y = 67 \text{ MPa}) , \end{aligned}$$

where s_1 , s_2 and s_3 are general principal stresses which are replaced by s_ϕ and s_z in this case. (The maximum shear stress is 2τ .)

The shear stress at the epoxy bond between the coil and support cylinder is another important mechanical parameter. We may roughly estimate this shear stress in the most pessimistic case by a simple

model. If the magnetic field at the coil end has only a radial component with a field intensity of 2 T, the estimated Lorentz force is

$$\begin{aligned} F_z &= B_r \times I \\ &= 16,000 \text{ N/m.} \end{aligned}$$

We assume that this axial force is supported by radial outer surface of the coil. Therefore, the approximate shear stress at the outer boundary is given by

$$\begin{aligned} s_{\text{shear}} &= F_z / W_z \\ &< 3.6 \text{ MPa,} \end{aligned}$$

where W_z is the axial width of the conductor. This number is at least four times smaller than the shear stress measured in various sample tests. The actual shear stress will be also smaller than that calculated with this simple method. A finite element analysis has been carried out to calculate this stress more precisely. Figure 5-8 shows the forces applied to the coil in the FEA and the resulting deformation, which gives us a strain level of less than 0.1%. Figure 5-9 shows the computed stress level. The maximum stress in the conductor is less than 52 MPa and the shear stress at the epoxy bond is 1 MPa or less. The conclusion is that the present solenoid design satisfies the mechanical design guidelines given above. A summary of the results of the stress analysis is given in Table 5-4.

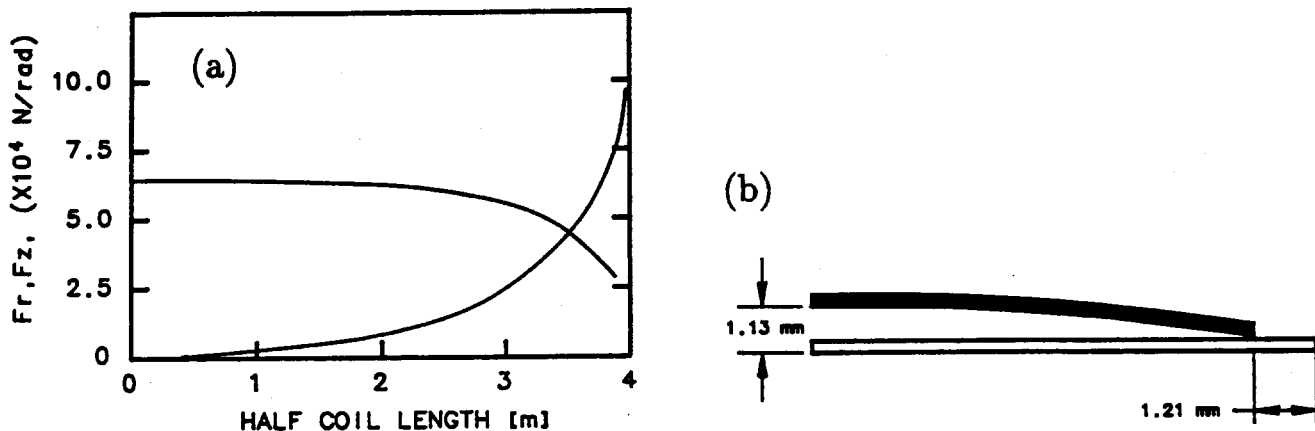


FIG. 5-8. (a) Radial and axial components of the force on the superconducting coil; (b) deformation of the coil under these forces.

Table 5-4
Mechanical stability of the SDC coil.

Element	Material	Yield stress	Anal. model	FEA
Conductor	Pure Al (3N8)	67 MPa	52 MPa	52 MPa
Cylinder	A-5083	169 MPa	52 MPa	< 52 MPa
Cylindrical bonding	Epoxy	(15 MPa)	< 4 MPa	< 1 MPa

In order to satisfy the design guideline of a maximum temperature less than 100 K after a quench—even without a dump resistor—(e.g., in case of some failure of a protection circuit), the stored energy must be

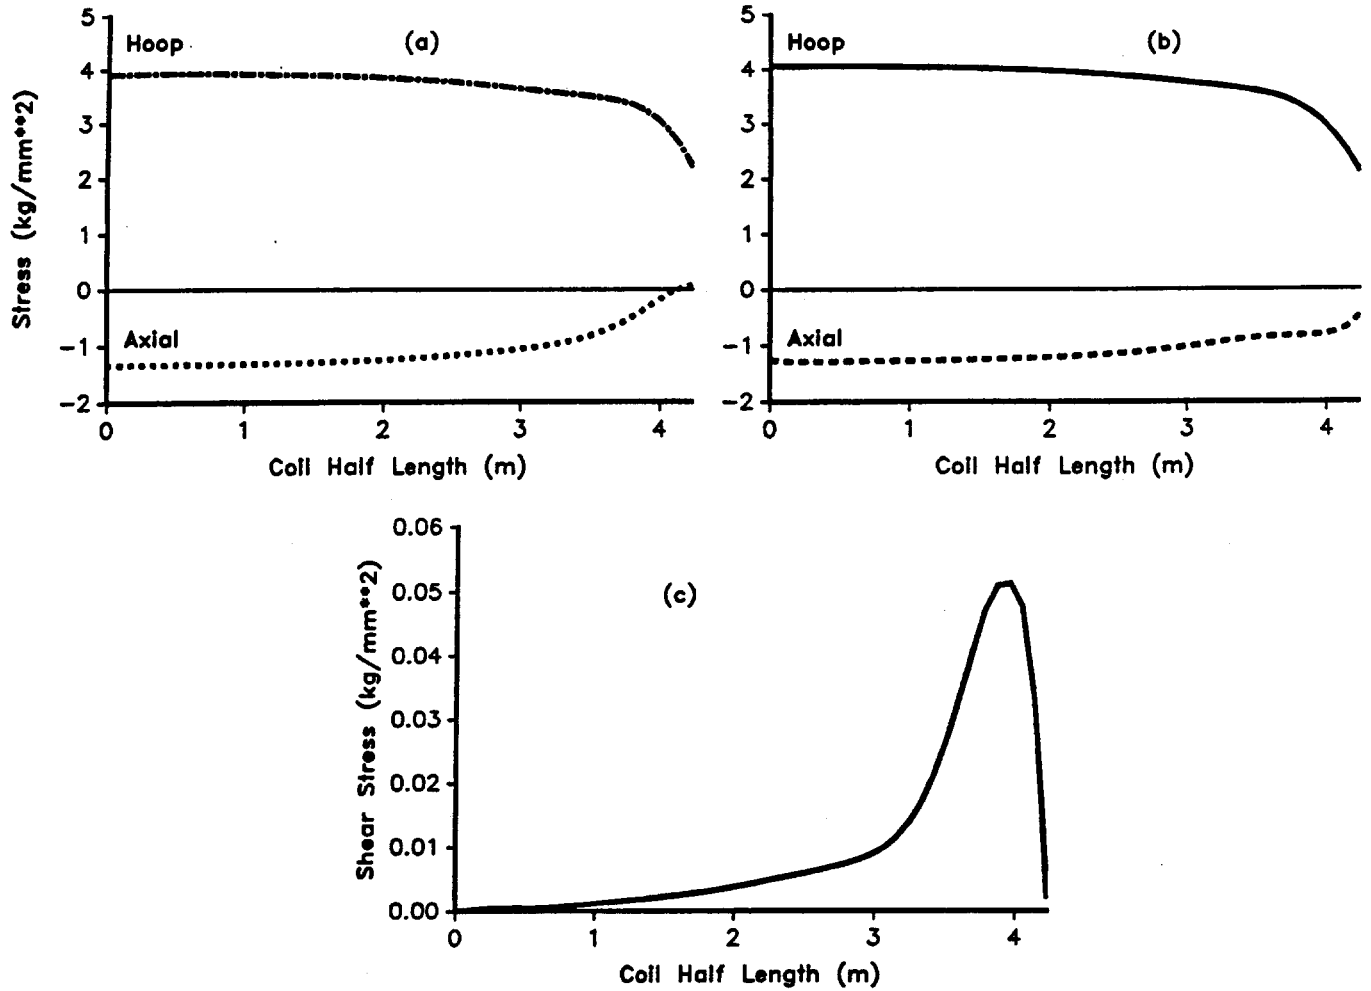


FIG. 5-9. Stresses as a function of axial position (a) on support cylinder, (b) on superconducting coil, and (c) on epoxy joint.

dumped as uniformly as possible into the coil. This is possible if the quench propagation time constant is much faster than the power decay time constant during the quench. A technique to increase the quench propagation velocity by using pure aluminum strips has been experimentally verified in the development of a thin superconducting solenoid for a balloon experiment [5,6]. If an adiabatic condition is assumed, the longitudinal quench propagation velocity $v(\phi)$ is

$$v_{\phi} = \frac{J}{\gamma_c} \frac{L_0 \theta_s}{(\theta_s - \theta_0)}.$$

where J is the current density, γ_c is the volumetric specific heat, L_0 is the Lorentz number, θ_s is the wavefront temperature and θ_0 is the initial operational temperature. The transverse quench velocity $v(z)$ may be written as

$$v(z) = \left(\frac{k_z}{k_{\phi}} \right)^{1/2} v_{\phi}$$

where k_z and k_{ϕ} are axial (transverse) and circumferential (longitudinal) thermal conductances in the coil. From these two equations we see that $v(z)$ may be enhanced by increasing the axial thermal conductance in the coil. Normally, the axial thermal conductance is suppressed by the turn-to-turn electrical insulation made of Kapton or glass tape. A pure aluminum strip pasted inside the coil surface serves to enhance

the effective thermal conductance in the axial direction by bypassing the axial electrical insulation. The maximum temperature after a quench has been calculated to be about 100 K in the SDC solenoid with an assumed 2 mm thick pure aluminum strip even in the most pessimistic condition that whole energy is dumped into the conductor only. Figure 5-10 shows the results of a computation of the temperature after a quench for dump resistors of 0.1 Ω , 0.05 Ω and 0 Ω , the last of which corresponds to the coil absorbing the entire stored energy. We expect another quench back effect from the outer support cylinder once the quench starts or a protection switch opens. This helps to distribute the energy absorbed more uniformly in the coil and lowers the final temperature.




FIG. 5-10. Quench behavior of the SDC solenoid with pure aluminum strips. Quench back in the outer support cylinder is not included in this calculation and the mass of the outer support cylinder is neglected.

5.3.3. Superconductor design and development

A high strength aluminum stabilized superconductor has been developed as part of the R&D program to meet the requirement of the SDC solenoid. Table 5-5 gives the parameters of the superconductor proposed for the SDC solenoid.

A superconducting cable is co-extruded in the center of the aluminum stabilizer. The cable, consisting of ten compacted strands, enables us to make cable lengths between joints much longer than that of the strands and also improves its critical current density because of its smaller filament size. Figure 5-11 shows the superconductor and the magnet characteristics in the SDC magnet design. The 8,000 A operating point will be 50% along the load line to the conductor capability at 4.2 K. The operating point corresponds to a temperature of 6.7 K and so the temperature margin is 2.5 K. The maximum field of 2.8 T includes a self field of 0.5 T at the coil end at 8,000 A. Based on recent experimental studies, the purity of the aluminum stabilizer has been optimized at 99.999% with 200 ppm zinc added. This improves the mechanical strength with little degradation of the residual resistivity ratio (RRR). Figure 5-12 shows the yield strength vs RRR for the aluminum with zinc. We can optimize these two parameters by using a combination of cold work and heat treatment effects during the cure process for the epoxy resin in the coil. A yield strength of 67 MPa is expected for the solenoid; our R&D work shows that an RRR of 500-600 can probably be achieved.

Stability of the superconductor against quench is one of the most important issues in the conductor design. It is parametrized by the minimum quench energy (MQE), which is a function of RRR. For the present conductor design, the MQE is expected from a simple one-dimensional model to be 0.27 J, which is sufficiently large. The RRR dependence of the maximum temperature after quench has been also verified with a computer simulation using the QUENCH code. It does not strongly depend on the RRR ratio in a region of 500-2000. Therefore, we conclude that an RRR of 500 is acceptable.

Table 5-5
Design parameters of the SDC superconductor.

Aluminum stabilized superconductor	
Superconductor material	Nb Ti/Cu
Stabilizer	Al (99.999% +200 ppm Zn)
Area ratio (Nb/Ti/Cu/Al)	1/1/29.8
Conductor size	$43.8 \times 4.37 \text{ mm}^2$
Superconductor strand	
Strand diameter	1.277 mm
Nb Ti filament diameter	20 μm
Number of filaments.	4100
Superconductor cable	
Overall size	$6.4 \times 2.5 \text{ mm}^2$
Number of cables	10 (2 \times 5)
Cabling pitch	50-60 mm
J_c in Nb Ti (at 5 T, 4.2 K)	2500 A/mm ²
Critical current (at 5 T, 4.2 K)	16,000 A
Characteristics of stabilizer	
RRR (Cu)	100
RRR (Al)	500
Yield Strength of Al (at 77 K)	67 MPa
Shear strength b/w Cu / Al	20 MPa
Conductor welding joint	
Joint resistance (at 4.2 K)	$< 10^{-9} \Omega$




FIG. 5-11. Short-sample characteristics of the SDC conductor and the operating load line.




FIG. 5-12. Residual resistivity ratio (RRR) as a function of yield strength in pure aluminum. The operating stress in the conductor is 52 MPa and the minimum acceptable RRR is 500.

5.3.4. Radiation shield

Two cylindrical radiation shields are placed inside and outside the coil, between the coil and the vacuum walls. The shields are made of A5083-H32 with a thickness of 2 mm. These two cylinders are jointed by end flanges to become one coaxial unit which is supported from the vacuum bulk heads through axial GFRP rods. The shield is broken electrically in the circumferential direction to eliminate the forces due to eddy currents induced by a fast discharge of the magnet. The path of the 15 mm inner diameter cooling pipe on the shield is similar to that on the outer support cylinder. It has 16 serpentine passes in series with a total pipe length of 320 m. The cooling pipe is designed in accordance with the ASME piping code.

5.3.5. Cryostat vacuum vessel

The vacuum vessel for the SDC solenoid consists of inner and outer coaxial shells with flat annular bulkheads on the ends. The cold mass is suspended inside the vacuum vessel by the support system. The superconducting leads and the cryogen pipes leave the vacuum vessel through a chimney at one end of the outer shell. The inner shell is 6 mm thick, fabricated of solid 5083 aluminum plates. The bulkheads are also 5083 aluminum and are 30 mm thick.

If the outer shell were made of solid aluminum plate, the radiation thickness of the solenoid would be greater than the requirement given in Table 5-1. An arrangement of more sophisticated materials must, therefore, be utilized. The predictability and reliability of such materials must be understood and the risk of using them found to be acceptable before selecting one for the outer shell. Radiation hardness is also very important in selection of the material and, therefore, no O-ring seal structure is being considered.

The requirements on the outer vacuum shell are summarized in Table 5-6.

Table 5-6
Outer vacuum shell requirements.

Vacuum load	1 atm radial and axial
Design standard	
Radial collapse pressure	> 2 atm
Axial buckling load	> 2 atm
Allowable stress	Based on ASME press. vessel code
Material	Aluminum alloy
Construction	Welded joints

One of the two techniques under consideration for the outer shell of the vacuum vessel is "isogrid"—integral, grid-stiffened shells and structures in which the grids form an equilateral triangle pattern. Isogrid has exceptional strength-to-weight and stiffness-to-weight characteristics. The repetitive pattern is also an important factor in minimizing costs by allowing effective use of numerically controlled (NC) manufacturing methods. The design of an isogrid vacuum shell consists essentially of solving simultaneous equations for general stability, local (rib and skin) stability, and stress loading with constraining equations (e.g., overall height). Metal isogrid shells are typically fabricated in steps: the grid pattern is first NC machined in flat plates, then the flat plates are formed by brake or roll forming into short cylinders. These cylinders are then assembled and welded to make up the shell.

A consulting firm has developed an isogrid vacuum shell for the SDC solenoid [7]. The goal of the design effort was to minimize the effective (weight-average) thickness of the isogrid under the requirements given in Table 5-6. The consultant used an outer radius of 2032 mm and a length of 9144 mm for a preliminary design. For comparison purposes, a 46 mm thick isogrid is given in Table 5-7.

Isogrid has never been used for such a large vacuum shell as that required for the SDC solenoid. Therefore, the outer shell of the vacuum vessel of the prototype magnet will be made of isogrid to demonstrate the feasibility of this technique.

Another option for the outer vacuum wall is "honeycomb"—a sandwich of two skins and a low-density hexagonal core. Honeycomb acts somewhat like an I-beam. The skins correspond to flanges and carry the tensile and compressive stresses, whereas the core corresponds to the web of an I-beam which carries the shear and helps to prevent buckling and wrinkling of the skins. The stiffness, or flexural rigidity, of a honeycomb panel is determined to first order by the thickness of the skins and their separation; the core provides no stiffness to the panel. Therefore a honeycomb is a very efficient type of structure with a higher stiffness-to-weight ratio than isogrid. In the past, aluminum honeycombs have been made by epoxy bonding the skins to the core. Bonded honeycomb has never been used for a shell as large as that required for the outer vacuum shell of the SDC solenoid and the long-term reliability of such a shell is unknown. Recently, large brazed aluminum honeycomb panels have been developed for industrial uses such as building siding and transportation vehicle walls [8]. Such a material has none of the disadvantages of bonded honeycomb and should be very reliable, since the vacuum seal and the joints between panels can be welded. To make sure that brazed honeycomb is a possibility for the SDC solenoid, an R&D effort is planned which will provide experience in the fabrication of curved honeycomb panels and their assembly into a welded shell. Table 5-7 compares a honeycomb shell with isogrid and solid shells. The decision on the material for the outer shell of the cryostat will be made after the R&D and the prototype programs have been completed.

Table 5-7
Comparison of solid, isogrid, and honeycomb outer vacuum shells.

	Solid	Isogrid	Brz. honeycomb
Aluminum alloy	5083	5083-H32	6951/4045-T6
Total thickness (mm)	27	46	46
Skin thickness (mm)	—	4.0	3.0 + 3.0
Skin layers	—	single	double
Node configuration	—	triangle	hexagon
Effective thickness (mm)	27	11	7
Weight reduction ratio	1	1/2.5	1/3.9
Radiation thickness (X_0)	0.303	0.123	0.079
Maximum size of plate (m \times m)	2 \times 6.4	2.2 \times 4.3	1.2 \times 4
Units to be welded	8 (= 2 \times 4)	12 (= 3 \times 4)	21 (= 3 \times 7)

5.3.6. Coil support

The purpose of the support system is to transmit both the cold mass and electromagnetic decentering loads from the 4.2 K superconducting coil package to the hadron calorimeter at 300 K. The cold mass of the outer support cylinder, conductor and insulation is approximately 20 metric tons. The magnetic forces are larger than the gravitational forces and thus dominate the design of the support structure.

The expected magnetic field distribution is such that if the geometric center of the coil were located at the magnetic center of the hadron calorimeter, the coil would be in unstable equilibrium with no net electromagnetic body forces acting upon it. However, if the coil were displaced from its equilibrium position, forces would act upon the coil in such a way as to increase the displacement.

For the purposes of designing the support system, it is assumed that the coil is 20 mm from the magnetic center of the hadron calorimeter in any direction. This choice was made based on the following considerations:

1. The coil can only be located within 5 mm of the geometric center of the hadron calorimeter due to tolerances in the calorimeter itself.
2. The scintillator slot pattern may cause the geometric center of the hadron calorimeter not to coincide with the magnetic center.
3. There may be variation in the magnetic material properties of the flux return.
4. The return paths contain several air gaps whose dimensions could change, especially the air gaps between the endcaps and the barrel portion of the calorimeter.
5. The space available to attach the coil to the barrel calorimeter and the limited access to this space makes attractive a design that does not require precision radial adjustment of the coil position to place it on true magnetic center. This adjustment is discussed below.

The magnitude of the calculated decentering forces for the baseline design with a 20 mm displacement are assumed to be 40 metric tons in the axial direction and 5 tons in the radial direction. However, a radial load of two g 's (40 tons) is used for designing the support system. This provides a margin for the expected shipping accelerations. The forces and force constant are given in Table 5-8. If the support system is to be stable and keep the deflections small, it must have a spring constant greater than the decentering force constant. This stiffness requirement is in addition to the load requirement.

Table 5-8
Design load of the coil support system.

Function		Upward	Downward	Left/right	F/B-ward
Global force					
Coil weight	(ton)	-20	20		
Applied load	(ton)	40	40	40	40
Design load	(ton)	20	60	40	40
Force constant	(t/mm)	2	2	2	2
Support stiffness	(t/mm)	>10	>10	>10	>10

The dedicated type of support system chosen consists of axial members to provide longitudinal stiffness and approximately tangential members to provide radial stiffness. The members connect between the outer support cylinder and the flat annular bulkhead of the vacuum vessel. A set of 8 or 12 axial, compression-tension members is located on the chimney end of the cryostat only while the tangential tension members are located at both ends. Both types of members are fabricated of epoxy-fiberglass (GFRP) composite and have design safety factors on the ultimate strength of at least four. The axial members are also designed for a buckling safety factor of 4. The GFRP members have thermal intercepts which operate near 80 K and at 4.5 K. Axial and radial differential contraction of the coil-outer support cylinder cold mass is accommodated by spherical bearing on both ends of each support member so that bending stress is not transmitted to the member. The metallic bearing and attachment block are designed for a safety factor of four.

The magnet cryostat is attached to the end modules of the barrel calorimeter as shown in Fig. 5-13. The axial attachment is at the chimney end only. Eight attachment points are spaced every 45°, corresponding with the 32 module segmentation of the barrel calorimeter. Each attachment allows 3 mm of radial movement, but restricts axial motion. The cryostat is radially connected to the calorimeter at both ends. The attachments have radial keys to provide for radial thermal contraction of the cryostat with respect to the calorimeter. The barrel calorimeter will be completely assembled prior to installation of the magnet cryostat. A fixture is used to install the cryostat into the barrel with the end opposite the chimney inserted first. The shims and keys are then installed.




FIG. 5-13. Cryostat attachment to the barrel hadron calorimeter, (a) section view, (b) enlarged end view.

5.3.7. Chimney

The cryogen lines and the superconducting leads on the solenoid magnet extend through the detector in a chimney to the service port outside the muon detector. The lines and leads are surrounded by a multi-layer insulated thermal shield and an outer vacuum jacket. The vacuum jacket has a section with a sliding sleeve for access to the field joints in the cryogen lines, the superconducting leads and the thermal shield. It also has a connector box to house the hermetic high voltage and instrumentation connectors. The 9.8 m long, cylindrical chimney runs vertically upward through the detector. It is attached to the top of the outer vacuum shell of the cryostat at the north end and to the bottom of the service port.

The stainless steel vacuum jacket is based on the ASME Pressure Vessel Code for full internal vacuum. It is 300–320 mm in diameter between the service port and an aluminum-to-stainless transition joint close to the magnet cryostat.

The chimney is supported on the muon iron. A stainless steel bellows is installed on the lower part of the vacuum jacket to accommodate relative motion between the muon iron, the chimney and the cryostat. This motion may arise from ambient temperature changes, cooling of the cryostat outer vacuum shell or chimney vacuum jacket by a thermal short or a cryogen leak, or from alignment adjustments. The cryogen lines, the thermal shield and the superconducting leads are designed to accommodate thermal contraction from room to operating temperature.

The annular space between the thermal shield and the vacuum jacket is a flow channel for pressure relief and pumping of the cryostat vacuum vessel. Pressure relief is needed to prevent an overpressure of the cryostat vacuum vessel by an accidental flow of cryogen from a ruptured or leaking line. Continuous vacuum pumping is required to keep the solenoid operating in the event of a small helium leak inside the cryostat vacuum vessel, chimney or service port. The relief device and vacuum pump-out are located on the service port.

Detector components are designed to accommodate the chimney. A notch, called the chimney wedge, will be provided in one or two of the calorimeter modules by removing two towers. A hole will be provided through the muon iron and chambers.

The connector box/field joint area is accessible only from one side of the chimney wedge. The field joints in the cryogen lines and the superconducting leads are arranged in the chimney such that adequate access is available during installation. Cables from the central tracker are routed around the open side of the chimney wedge to provide access to the connector box and the removable section of vacuum jacket.

5.3.8. Thermal design and cooldown characteristics

The coil-and-support cylinder assembly is cooled by the forced flow of two phase helium inside a single pass cooling tube. The CDF and TOPAZ thin solenoids are cooled in this manner. Their operating performance has demonstrated that indirect cooling is effective and reliable. This method was chosen because it is simple and requires less radial space.

The thermal radiation shield will be cooled with either gaseous helium from an intermediate stage in the refrigerator or with forced flow liquid nitrogen. Having only one cryogen in the system should simplify operation and maintenance.

Estimates of the normal operating, steady state heat loads including a contingency of 20% for the magnet, chimney and service port, are given in Table 5-9. Cooling design parameters for cooldown, steady state, excitation and quench recovery are given in Tables 5-10, 5-11, and 5-12.

Table 5-9
An estimate of steady state thermal loads for the SDC solenoid.

Component	300 to 77 K	77 to 4.2 K	300 to 4.2 K
Thermal radiation	300	30	—
Conduction			
Coil support rods ($n = 24-36$)	24	2.5	—
Shield support rods ($n = 12$)	3	—	—
Chimney and service port	36	4.5	—
Current leads (8 kA pair)	—	—	30
Total thermal load	363 W	37 W	30 L/hr

Table 5-10
Cooling design parameters in the initial cool down.

Boundary conditions	
Coil cold mass (aluminium)	20 tons
Radiation shield mass (aluminium)	1 ton
Temperature difference in coil during cool down	<50 K
Coil cooling	
Initial cooling speed	1 K/hr
GHe mass flow required	21 g/s
Initial cooling power	4.4 kW
Initial inlet pressure	0.7 MPa A
Pressure drop in cooling pass (25 mm $\phi \times 250$ m)	0.2 MPa
Shield cooling	
Initial cooling speed	1 K/hr
GHe mass flow required	1.2 g/s
Initial cooling power	250 W
Initial inlet pressure	0.7 MPa A
Pressure drop in cooling pass (15 mm $\phi \times 400$ m)	<0.035 MPa

Table 5-11
Cooling design parameters in steady state operation.

Coil			
Typical coil temperature			4.4 K
Quality of two phase helium	Inlet		0.8
	Outlet		0.5
He mass flow required			7 g/s
Pressure drop in cooling pass			small
Current leads mass flow			1-2 g/s
Radiation shield			
Max. shield temperature			< 70 K
Inlet He gas temperature			60 K
Helium gas mass flow required			15 g/s
Heat exchange efficiency			0.5
Pressure drop in cooling pass			small

Table 5-12
Cooling parameters in excitation and after quench.

Excitation of the coil			
Static heat-in-leakage			37 W
Eddy current loss in support cylinder	(@ 1 mT/s)		88 W
Quality of two phase helium	Inlet		0.8
	Outlet		0.5
He mass flow into cooling pass required			23 g/s
He mass flow into current leads	(2 × 8,000 A)		1 g/s
Quench recovery in coil			
Energy dumped into coil			88 MJ
Recovery time assumed			4 hr
Cooling power required			6.2 kW
Cooling efficiency assumed			0.8
Liquid (two phase) helium flow required			28 g/s (830 L/hr)
Total liquid helium for recovering			3,300 L

5.3.9. Diagnostics and instrumentation

A substantial amount of instrumentation is required to monitor the cooldown and steady-state operation of the magnet and to provide diagnostic data. The monitoring and diagnosis functions include:

- Maintaining the appropriate cool-down rate, allowable temperature distribution, mechanical stress and thermal contraction.
- Monitoring the magnet during excitation and operation by measuring the temperature of the cold mass, thermal shields and support intercepts; the stress in the outer support cylinder and supports; the magnet current and voltage; and the magnetic field.
- Monitoring the magnet during and after a quench to verify the safety of the system. Current, voltage, temperatures, pressures and stresses will be measured as a function of time.

We will need to provide a sophisticated monitoring system with a link to the refrigerator control

Table 5-13
Diagnostic instrumentation in the SDC solenoid.

Element	Coil	R/S	Support	V/V	Chimn.
Voltage taps (short pair tap)	13				4
Quench heater	3				
Temperature sensors					
Pt-Co (300 K-4.2 K)	24	24	4	4	6
CGR (high resol. @4.2 K)	17				
Strain gauge (pair of ϕ , z)	12	8	36	16	4
Position sensor (r , ϕ)	8				8
(z)	4				4
Pressure	3				
Total	83	32	40	20	26

system. Table 5-13 gives a list of diagnostics and instrumentation required.

5.3.10. Magnet assembly and initial test

The SDC solenoid coil will be wound with an inside-out winding technique in a vertical configuration. The magnet will be assembled vertically to minimize fixtures and tooling. The magnet will be tested in air after the assembly before shipping to the SSC Laboratory.

5.4. Transportation and installation

The present transportation plan is to air ship the magnet from Japan to the Dallas/Fort Worth airport to minimize mechanical shock and handling. Heavylift Cargo Airlines can carry the 4.1 m \times 9 m cryostat on their AN-124 aircraft. The magnet will be moved by road from the airport to the SSC site.

The magnet will be installed into the SDC detector after the assembly of the barrel calorimeter structure. Special fixtures will be developed for this purpose. The field joint between the cryostat and the chimney will be made as the last step in the installation sequence.

5.5. Cryogenics

5.5.1. General requirements

General requirements for the SDC system, summarized in Table 5-14, have been set based on the superconducting solenoid magnet system cooling requirements described above, the possible use of a VLPC (Visible Light Photon Counter) system described in the previous chapter, and other system heat loads, such as the transfer line and the liquid helium storage dewar. The capacity is sufficient for the nonsimultaneous cooldown and steady-state operation of both the solenoid and VLPC systems. Cooldown of the superconducting solenoid magnet after a quench is effected by using liquid helium from the storage dewar. Capacity of the refrigerator/liquifier is not sufficient to reliquefy the helium gas generated while cooling down the magnet after a quench. The gas is stored and can be reliquefied in about 12 hr. Liquid helium inventory and helium gas storage capacity is sufficient for recovering the solenoid from two quenches per day. The VLPC system remains in operation (not cooldown mode) while the quench-recovery gas is reliquefied. Approximate heat loads other than the superconducting solenoid magnet are: transfer line, expansion boxes and U-tubes, 50 W; control dewar, 13 W; 10,000 liter liquid helium storage dewar, 4 W; equivalent refrigeration load for the vapor cooled lead liquefaction load of 1 g/s, 100 W; and the VLPC

Table 5-14
General requirements on the cryogenics for the SDC solenoid.

Typical refrigeration capacity	1,500 W at 4.4 K
Liquid helium transfer rate	750 L/hr (after quench)
Liquid helium storage capacity	5,000–10,000 L
Cold gas helium mass flow at 60 K	12 g/s




FIG. 5-14. Flow schematic of the cryogenic system for the SDC solenoid.

system, to be determined. The above comments regarding the VLPC system are not final because the cooldown and steady-state loads for it are not well defined at this time.

5.5.2. SDC cryogenic system

The SDC cryogenic system consists of the helium refrigerator/liquefier, liquid and gas transfer lines, liquid and gas storage, and possibly a nitrogen system. It is designed for the loads discussed in the previous subsection. A schematic diagram of the SDC cryogenic system is shown in Fig. 5-14.

The refrigerator/liquefier is housed in a surface utility building. Major components in it include: the compressors, oil removal system, coldbox and control system. The liquid helium storage dewar and helium gas storage tanks for this system are located outdoors adjacent to the building. All the helium supplied by this system, except for normal leakage and necessary venting, is recovered and reliquefied.

The 10-inch transfer line, described in Section 12.5.5, terminates in the control dewar. The control dewar is the interface between the transfer line and the superconducting solenoid and VLPC systems. It is located next to the solenoid service port at the upper end of the chimney and is connected to it with U-tubes. The control dewar is installed on the end of the transfer line with a vacuum break.

The control dewar contains a helium subcooler which is used to remove transfer line heat leak, valves for controlling the flow of liquid helium to the cryostat, a cooldown by-pass valve and instrumentation sensors. It has bayonets for U-tubes which connect it to the service port, and to the liquid helium circuit

of the VLPCs. It also has trapped volume relief valves on the cryogen lines and hermetic instrumentation connectors. A thermal shield and intercepts reduce the heat flux to the liquid helium components. The vacuum vessel has access panels and reasonable interior space for doing repair work. The cryogenic pressure vessels and piping are designed in accordance with the ASME Piping and Pressure Vessel Codes.

The control dewar is supported from the muon iron. The section of transfer line between the control dewar and an expansion box mounted on the west wall of the detector hall has a bellows to accommodate relative motion between the control dewar and the transfer line to the surface.

The nitrogen system supplies both liquid and gaseous nitrogen to systems and components in the SDC experiment. Possible liquid nitrogen uses include cooling the transfer line thermal shield, cooling the thermal shield and thermal intercepts in the superconducting solenoid magnet, cooling the high pressure helium feed stream in the refrigerator/liquefier and auxiliary uses. All of the nitrogen supplied by this system vents to the atmosphere after use. Major components include the storage dewar and subcooler. The storage dewar is located on the surface in the same areas as the helium dewar and tanks. The subcooler is located either adjacent to or inside the control dewar. The effluent nitrogen stream from the superconducting solenoid magnet enters the shell side of the subcooler and boils to subcool the supply stream in the tube side. Only nitrogen gas vents to the surface from the shell side of the subcooler. Liquid nitrogen is intentionally kept out of the vent line to minimize the shell side operating temperature. An internal heater is used to boil excess liquid nitrogen in the shell side to ensure the vent line stays free of liquid.

5.6. Electrical system

5.6.1. General features

The electrical power system is required to satisfy the the following functions:

- To energize the solenoid up to a rated current within one hour.
- To hold the magnet at operational current for long periods—typically days.
- To detect a quench and to switch off the secondary circuit and to extract the stored energy from the magnet.
- To discharge the magnet slowly in case of cryogenic failures or similar interruptions. The discharge rate should be slow enough to eliminate a quench back due to eddy current loss in the support cylinder.

The electrical system consists of a power supply, voltage filter, fast and slow discharge resistors, and high-current DC switches as shown in Fig. 5-15. The power supply and filter are located in the service building on the surface; discharge resistors and the DC switches are located in the detector hall. Water-cooled bus runs from the service building to the detector hall. Table 5-15 gives general requirements and operational condition for the electrical power system.

5.6.2. Power supply

The power supply consists of twelve phase thyristor regulator, filter and control circuit for constant current/voltage operation. The thyristor regulation may be allowed with filtering system. A passive filter system reduces voltage ripple under the level of a few hundred volts. The current stability is required to be 0.01% at the nominal operational current with an inductive load of 4.6 H and a resistive load of 1.2 m Ω through a current bus-bar to the solenoid magnet. A voltage stability of 0.1% may be required to energize the magnet in voltage constant mode without excessive overshooting in an optional operation mode. The bus bar system runs with a total length of about 200 m from the power supply to the magnet. The power supply system including the bus bar system requires de-ionized water with a bulk resistivity of 1 M Ω m or more for cooling the system. The size of the power supply is expected to be 2 \times 10 \times 2.2 m³.

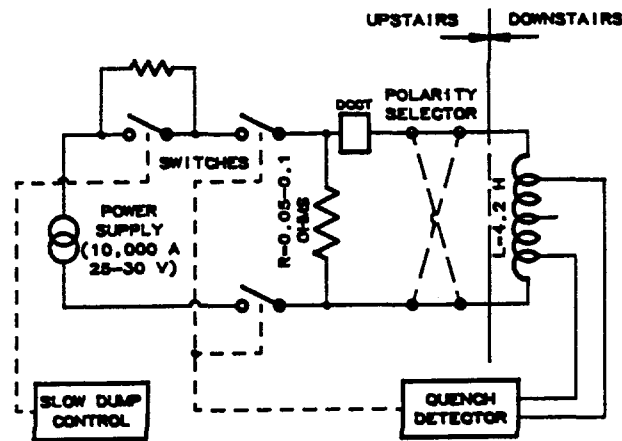


FIG. 5-15. Schematic of DC circuit for SDC solenoid.

Table 5-15

General requirements for the SDC electric power system.

Maximum current	10,000 A
Nominal operation current	8,000 A
Regulation	$< \pm 1 \times 10^{-4}$
Inductive voltage (@1 mT/s)	± 17 V
Bus-bar voltage drop	10 V
Ramp rate	1-10 A/s
Fast discharge time const.	45 s
Slow discharge time const.	500 s
DC switching time	< 1 s

5.6.3. Discharge resistors and switches

Three modes of switching and discharging are provided to extract the electromagnetic stored energy from the magnet. A fast discharge system will quickly switch off the circuit and extract the energy to a dump resistor. The system must operate during a power failure and, therefore, the switch should be of the power-to-close type and the dump resistor should be an adiabatic type, cooled by natural convection only. Redundancy will be provided to the fast switching system and the quench detectors because they are important elements of the quench protection system. A slow discharge system will discharge the magnet sufficiently slowly that eddy current in the outer support cylinder will not initiate a quench. It will also be possible to discharge the magnet very slowly by dialing down the current set point.

5.7. Prototype R&D

5.7.1. Prototype coil

The development of a full diameter, one quarter length prototype magnet has been underway since 1991. The purposes of the prototype development are:

- to develop an aluminum stabilized superconductor of high yield strength and RRR;
- to become skilled in the fabrication of a full diameter coil using inner winding techniques;
- to demonstrate that a coil with E/M of about 8 kJ/kg can be quenched safely and without damage or reduced performance; and
- to apply electromagnetic loads to the conductor and outer support cylinder equal to those expected in the detector magnet and to verify the performance under these loads. This can be achieved by operating the prototype to 10–12 kA.

The basic development of the prototype superconductor has been completed and the conductor for the prototype is being fabricated. Table 5-16 gives the main parameters of the prototype.

Table 5-16
Design parameters of the prototype R&D solenoid.

Dimensions		
Cryostat	Inner radius	1.70 m
	Outer radius	2.06 m
	Half length	1.17 m
Coil	Effective radius	1.85 m
	Half length	0.95 m
Conductor	Thickness	44 mm
Outer cylinder	Thickness	33 mm
Transparency	Radiation thickness	$1.23 X_0$
	Interaction length	$0.26 \lambda_0$
Electrical parameters		
Central field		1.54
Nominal current		11,250 A
Inductance		0.68 H
Stored energy		48 MJ
E/M		10 kJ/kg
Mechanical parameters		
Effective cold mass		4.5 tons
Total weight		8 tons
Radial mag. pressure (@z = 0)		1.73 MPa
Axial compressive force		16.7 MN
Maximum hoop stress		43 MPa
Maximum axial stress		-19 MPa
Maximum shear stress		62 MPa
Peak field in coil		3.8 T
Load line ratio		70%

5.7.2. Prototype vacuum chamber

The vacuum vessel for the prototype magnet cold-mass consists of an isogrid outer shell, a solid inner shell and two flat, annular bulkheads, all made of aluminum. The vessel is designed with O-rings and bolts so that both bulkheads may be removed for inspection, modifications, or repairs.

The isogrid pattern designed for the outer shell of the production solenoid will be used for the prototype vessel so that fabrication experience relevant to the detector magnet vessel will be gained. A flanged chimney nozzle 323.8 mm (12.75") in diameter penetrates the shell near one end; the grid is reinforced around the nozzle for the vacuum punch load. The isogrid pattern will be NC machined into three flat plates of 5083 aluminum alloy. The plates will be formed into cylindrical sections with the pattern on the inside and welded together to make the shell. The inner vacuum shell is made of 5083 aluminum alloy and has a wall thickness of 6.4 mm (0.25"). Its collapse pressure is 41 kPa (6 psi).

The bulkheads will be fabricated from 32.0 mm (1.25") aluminum plate, with O-ring grooves machined near the inner and outer perimeter. Both the radial and axial supports will be attached to the bulkheads. Figure 5-16 shows the cross sections of the prototype magnet.

5.8. Safety and protection

A Conceptual Hazard Assessment Report (CHAR) of the superconducting solenoid system is part of the draft Conceptual Safety Analysis Report prepared in conjunction with the Design Report. The CHAR is a qualitative study of the potential hazards associated with the solenoid and the methods that will be used to mitigate those hazards.

5.9. Schedule

The SDC solenoid will require a seven-year-long program to be completed. The prototype phase lasts three years; the production magnet fabrication should take three years, including the initial test in air. The final phase for transportation, installation and testing in the SDC detector needs two years, of which some overlaps with the previous phase. The overall schedule for the SDC solenoid development is given in Table 5-17.

References:

1. SDC Solenoid Design Notes, Fermilab. (1990-1991)
2. T. Kirk *et al.*, Report of the Magnet Task Force, Solenoidal Detector Collaboration, (1991)
3. A. Yamamoto *et al.*, *Proc. of the Int. Workshop on Solenoidal Detector for the SSC*. KEK, Tsukuba (1991) p. 141.
4. I. Inoue *et al.*, to be published in the *Proc. of IISSC* (1992)
5. A. Yamamoto *et al.*, *IEEE Trans. MAG*, 24 (1988) p. 1421.
6. Y. Makida *et al.*, *Adv. Cryog. Eng.* 37, 401 (1992).
7. L. W. Swenson, Fermilab SDC Solenoid Design Note No. 154 (1991)
8. K. Okuto *et al.*, Sumitomo Light Metal Technical Report Vol. 32, #4 (1991).




FIG. 5-16. SDC prototype solenoid. Dimensions are given in Table 5-16.

Table 5-17

Overall schedule for the prototype solenoid and the production solenoid. The Japanese fiscal year begins April 1 of the year indicated.

JFY1991	Prototype Magnet Development —Superconductor fabrication —Winding machine development —Outer support cylinder fabrication —Isogrid vacuum wall development
JFY1992	—Coil winding —Cryostat element fabrication
JFY1993	—Assembly of the magnet —Cool-down and excitation in air
JFY1994	Production Magnet Fabrication —Superconductor fabrication —Cryostat element fabrication
JFY1995	—Coil winding —Magnet assembly
JFY1996	—Magnet assembly continued —Cool-down and excitation test in air —Transportation to SSCL
JFY1997	—Cool-down and excitation in iron —Field mapping

6. Calorimetry

6.1. Introduction

The basic functions of the SDC calorimeter systems are to identify electrons and photons and measure their energies (in conjunction with the tracking system), to measure the energies and directions of jets, and to provide hermetic coverage for measurements of missing transverse energy. The physics requirements that guide the design of the calorimetry are reviewed in Section 6.2. The calorimeter system is described in Section 6.3. The central pseudorapidity range ($|\eta| < 3$) is covered by barrel, endcap and plug calorimeters surrounding the solenoid and tracking volume. These calorimeter elements provide the detailed electromagnetic shower information required for electron and photon identification and energy measurements and determine jet energies and directions. Hermetic coverage is completed by the forward calorimeter system, which covers $3 < |\eta| < 6$.

In the Letter of Intent, a liquid argon option and a scintillation option were described for the central region. [1] We have chosen scintillation calorimetry with lead and iron absorbers for the electromagnetic and hadronic sections, respectively. The scintillation detection elements are divided into tiles, each tile being read out by a waveshifting fiber. The optical readout design is described in Section 6.4. An important design consideration is to maintain high light yield. Masking, which reduces average light yield, will be needed to provide longitudinal and transverse uniformity. The large number of tiles and fibers requires efficient fabrication techniques and quality control. The sensitivity of the optical system to the radiation environment resulting from the high luminosity of the SSC has motivated much study of both components and test calorimeter modules (Section 6.5). The challenge of finding an affordable mechanical design which does not compromise the physics measurements has been met (Section 6.6).

The central calorimeter contains a shower maximum detector (SMD) located near the electromagnetic shower maximum, which measures transverse shower profiles. The parameter choices of the SMD are based on considerations of electron and photon identification. The first scintillator tile in each calorimeter tower is separately read out (massless gap) to aid in photon identification and to correct for electron energy lost in the coil. These elements are described in Section 6.7.

Systems for monitoring response in the tile/fiber calorimeter (Section 6.8) must serve both the function of quality control in establishing uniformity during assembly, and establishing and monitoring response levels in operation. Extensive radioactive source systems are used throughout the central calorimeter. A laser light flasher system is used to monitor photodetector gain and timing. These systems, together with test beam data and *in situ* analysis of selected data samples, provide the absolute and relative calibrations of response.

Section 6.9 covers the broad spectrum of issues for the central calorimeter that remain to be resolved through engineering, research and development and full-scale prototype fabrication.

We are considering two options for the forward calorimetry: high pressure gas ionization readout or liquid scintillator in small tubes. Design of the two forward calorimeter options has begun, but a substantial R&D effort is required to select the final technology (Section 6.10).

Finally, in Section 6.11, we summarize the results from two recent test beam programs at Fermilab. Reconfigurable scintillators and radiators were used to study both principles of calorimetry and the response of candidate configurations for the central calorimeter. Prototypes of barrel electromagnetic modules, the shower maximum detector and hadron sections were also studied.

6.2. Physics requirements

6.2.1. Calorimeter environment

The significant challenges posed by the environment in which the calorimeter must function are:

1. The calorimeter is located outside of the solenoid magnet and the tracking system. The presence of even a thin solenoid in front of the calorimeter degrades the precision that can be achieved in the electromagnetic shower measurement, although much of the degradation can be recovered using the "massless gap" technique to sample showers initiated in the solenoid.
2. Significant magnetic fields permeate parts of the calorimeter. The calorimeter endcaps must withstand the full solenoid field and carry it to the return yoke. The response of the calorimeter and its readout in this field must be understood.
3. The maximum ionizing dose in the calorimeter occurs at the maximum of the electromagnetic showers. Radiation levels are modest in the barrel region, but reach extreme values in the forward calorimeter. The calorimeter performance requirements must be met, taking into account the impact of these radiation levels.
4. The 16 ns time interval between SSC beam crossings sets the scale for the time resolution needed in the calorimeter. The overall event rate of 100 MHz at design luminosity may rise by another factor of 10 over the life of the experiment. The calorimeter must help distinguish the rare interesting events from this background and must have the granularity and time resolution to suppress multi-event pileup.
5. There is non-negligible material in the tracking chambers between the beam collision point and the calorimeter face. The design must take into account the impact of this material on the calorimeter performance for electron identification.
6. The outer radius of the calorimeter determines the inner radius of the muon system. Thus, the size and cost of the muon system (and, of course, the calorimeter itself) depend on the calorimeter depth.

6.2.2. General calorimeter requirements

The physics requirements that influence calorimeter parameters are briefly discussed below and are summarized in Tables 6-1 to 6-3. Additional details can be found in Chapter 3.

Solid angle coverage of the EM calorimeter

Analyses of the physics processes $H \rightarrow 4e^\pm$ and $H \rightarrow 2e^\pm 2\mu^\pm$ guide the choice of rapidity coverage of the electromagnetic calorimeter. Typically, decreasing the coverage from $\eta = 2.5$ to $\eta = 1.5$ reduces the acceptance by a factor of 2.5 whereas increasing the coverage to $\eta = 3.0$ increases the acceptance by 20%. For example, the acceptance as a function of the pseudorapidity of the most-forward decay lepton for a 400 GeV Higgs is shown in Fig. 6-1. In order to retain high efficiency, the EM calorimeter is required to identify and reconstruct electrons within $|\eta| \leq 2.5$. The total coverage of the EM calorimeter must then extend to $|\eta| \leq 3$ to be able to establish isolation for these electrons. It has been shown that this coverage ensures high efficiency for a wide variety of high p_t processes. [2]

EM losses within fiducial volume

High electron-identification efficiency is essential to ensure sufficient rate for identifying rare processes such as the four-electron final state of the Higgs boson. In addition to losses due to expected inefficiencies in E/p matching, track-shower matching, isolation, or hadronic versus electromagnetic energy cuts, electrons can be misidentified because of imperfections in the calorimeter response. These imperfections include cracks and dead channels, non-gaussian tails in the energy measurement, and nonuniformities in response. The total loss from calorimeter imperfections should be no greater than 5% per electron with $E_t > 10$ GeV in order to retain an overall electron efficiency greater than 85%.




FIG. 6-1. Acceptance for 400 GeV Higgs particles in the decay mode $H \rightarrow ZZ \rightarrow 4\ell$ as a function of $|\eta|$ of the most forward lepton, where all the leptons are required to have $p_t > 10$ GeV/c (solid) or $p_t > 20$ GeV/c (dashed).

Solid angle coverage of the hadron calorimeter

Missing transverse energy is a signature of several physics processes, and calorimeter coverage over a large solid angle is essential to the precise measurement of missing transverse energy. One example is the production of gluinos, the supersymmetric partners of gluons (Section 3.4). A gluino can decay to jets plus the lightest SUSY particle, which, since it does not interact in the detector, produces substantial missing- E_t as shown in Fig. 3-71. This figure also shows the four-jet background for various ranges of calorimeter coverage. For a 300 GeV \tilde{g} signal to be visible, the calorimeter must be able to identify and reconstruct jets with axes within $|\eta| \leq 5.0$. The active area must be larger, approximately $|\eta| \leq 6$, in order to contain these jets. Similar solid angle coverage is required to tag the forward jets that are characteristic of the production of a heavy Higgs through WW/ZZ fusion (Table 3-9).

Energy losses within the fiducial volume of the complete calorimeter

Small cracks and regions of reduced response are somewhat less important for jets than for electrons. However, significant missing- E_t can be produced if, for example, the leading fragment of a jet or the hard photon in a γ -jet event is not observed. A cut requiring that the missing- E_t not be collinear with a jet will suppress events with the former problem, but not the latter. Preliminary calculations indicate that losses at the level of 1% produce a background that is comparable to the rate of $Z + \text{jet}$ events, where $Z \rightarrow \nu\bar{\nu}$. Gaps in the complete calorimeter (EM+HAC) must not exceed this level.

EM energy resolution

Two examples of physics analyses in which the resolution of the EM calorimeter is an important parameter are the measurement of the width of a heavy Z' boson, and the observation of a light Higgs boson through its decay to two photons. As discussed in Section 3.2.1, the signal-to-noise for the latter process is substantially improved by searching for its production in association with either a W boson or a $t\bar{t}$ pair. Identification of a significant signal requires good energy resolution with both stochastic and constant terms being relatively small. This study indicates that a stochastic resolution $\leq 15\%/\sqrt{E_t}$ and a constant term $\leq 1\%$ are highly desirable. Similarly, the measurement of the width of a Z' of mass 1 TeV requires a constant term of $\leq 1\%$ (see Section 3.5). This limit on the constant term implicitly sets limits on uniformity, stability, calibration, nonlinearity, leakage, radiation damage, and all other effects that contribute.

EM transverse segmentation

One of the characteristic features of an EM shower, which plays a significant role in electron and photon identification, is its small transverse size. The EM transverse segmentation should reflect this feature. Approximately 99% of the shower energy is deposited within 3 Moliere radii of the shower axis, corresponding to a radius of 0.05 in $\Delta\eta \times \Delta\phi$ at $\eta = 0$. Cells of size 0.05×0.05 , combined in a 2×2 array, form the minimal area for precision reconstruction of the shower energy. Such an array contains the energy, and does not suffer from edge effects. If the cell size were increased beyond the value of 0.05, performance would be degraded for the following reasons.

First, accurate calorimeter energy measurements at the SSC require minimizing the energy collection region in order to minimize pileup effects. Studies in Sections 3.2.1 and 3.2.2 indicate that the impact of pileup at higher luminosity is noticeably worse for 0.1 EM segmentation, and unacceptable for 0.2.

Second, high- p_t Z 's from the decay of a heavy Higgs produce di-electrons with small opening angles (see Fig. 3-19). Good efficiency for an 800 GeV Higgs requires reconstruction of electrons separated by $\Delta R \geq 0.3$. If we construct 2×2 cell arrays to measure the energy precisely and then demand that two arrays have a separation of at least two cells, this requirement implies segmentation at the level of 0.05×0.05 or better.

Hadronic energy resolution and transverse segmentation

The program most demanding of good hadronic resolution and segmentation would be the detection of narrow states decaying into pairs of jets. The dijet mass resolution receives contributions from physics effects such as fragmentation, as well as detector effects such as angular and energy resolution. When the jet p_t is small, mass resolution is dominated by physics effects. High- p_t jets may arise either from the decays of boosted light objects or from decays of heavy objects. For the boosted case, angular resolution plays a more important role than energy resolution. Only in the case of back-to-back high- p_t jets arising from the decay of a heavy object are the physics and angular effects suppressed to the point where the energy resolution plays a significant role.

We have studied the influence of hadron calorimeter transverse segmentation on the dijet mass resolution for hadronic decays of boosted W 's and Z 's. These might arise from the Higgs decay to WW or ZZ , or from $t\bar{t}$ events. Segmentation coarser than $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ significantly degrades the mass resolution, particularly for W/Z $p_t > 500$ GeV/c (see Table 3-7 and Fig. 3-30). The influence of the hadron calorimeter energy resolution was also studied, and Table 3-7 indicates that it has relatively little effect on high- p_t W/Z resolutions.

The only physics process in which the hadronic energy resolution is expected to have an important effect is the detection of a heavy and narrow object, such as a heavy Z' , decaying into two jets. Unless the production rate were relatively large, such a signature would be difficult to see above the QCD background. Given this paucity of strong physics constraints, we have chosen a moderate performance goal of $70\%/\sqrt{E_t} \oplus 6\%$ for single pions. For jets, the corresponding constant term in the resolution would be 2% (Table 6-8).

EM linearity

Accurate mass measurements of particles through their decays to electrons or photons rely on the linearity of the calorimeter; two examples are $H \rightarrow 4e^\pm$ and $Z' \rightarrow e^+e^-$. We therefore require the EM nonlinearity after corrections to be less than 1% for $E_t > 10$ GeV. The "massless gap" technique, described in the Introduction, will be used to correct for energy lost in the solenoid. A residual nonlinear response will degrade the mass resolution and bias the mean value.

Hadron calorimeter linearity

One signal for quark substructure would be an observed excess of jets at high E_t over the expected QCD rate (Section 3.6). The rate of these jets, however, is extremely small compared to the QCD rate at lower E_t . Errors in the jet linearity or in the calibration of the calorimeter can reduce the significance of an observed signal or artificially create a signal by systematically mismeasuring lower E_t jets. Studies presented in the EoI indicate that a residual nonlinearity (i.e., after linearity correction) for jets of 1% per TeV above 2 TeV will produce high- E_t backgrounds at a rate comparable to a compositeness signal from a 25 TeV contact term. Simulations of the jet response indicate that an error in understanding linearity of single-pion response at the level of 5% per TeV above 100 GeV produces a jet linearity error of approximately this magnitude (Fig. 3-80).

Dynamic range

The dynamic range required per EM or hadronic tower is approximately 20 MeV to 4 TeV in transverse energy. The lower limit allows an adequate measurement of 1% lateral leakage of 10 GeV electrons from the EM calorimeter and permits an isolation cut level as low as 100 MeV in the hadron calorimeter. The lower limit also enables muons to be identified at all angles in the hadron calorimeter, both for calibration purposes and as an additional aid in muon identification. The equivalent transverse energy deposited by a muon in a hadron compartment decreases from ~ 2 GeV at $\eta = 0$ to ~ 0.3 GeV at $\eta = 3$.

The upper end to the dynamic range is determined by the maximum energy per channel deposited by high- E_t jets. A study has been performed to isolate those physics processes which are likely to produce the largest energy depositions in single cells. For this purpose, the signal from a 4 TeV Z' with Standard Model couplings and the signal from high- p_t jet production ($p_t > 5$ TeV/c) were generated. The simulation used the PYTHIA generator, followed by the shower parametrizations described in Section 3.1.1. The rate of events depositing more than 4 TeV of transverse energy in a single EM or hadronic tower is on the order of several per SSC year (Fig. 6-2). Energy deposits larger than this amount will produce nonlinearities in jet response that must be understood to the level discussed earlier. The 4 TeV limit provides sufficient range to accurately measure the energy deposited by an electron from the decay of a 4 TeV Z' (Fig. 6-2a).

EM depth

Leakage from the EM calorimeter contributes to the constant term in the EM energy resolution. This contribution can be reduced by using the signal from the relevant hadronic compartment to make a leakage correction. To keep the constant-term contribution from EM leakage below 0.3%, an EM calorimeter depth of $22 X_0$ is adequate with a hadronic calorimeter having 2.5 cm iron absorber plates, while a depth of $25 X_0$ is required if the backing calorimeter has 5 cm plates [3].

Total calorimeter depth

The total calorimeter depth is required to be 10λ at $\eta = 0$ to provide sufficient containment for high energy jets [4]. Figure 6-3 shows the mass resolution, σ_M/M , for 10 TeV dijets vs total calorimeter depth for three different sets of test beam data that were used to parametrize the response function. Although the data are not entirely consistent, a 10λ depth appears to just provide sufficient containment. A thinner calorimeter would not only degrade the resolution, but also the efficiency, by increasing non-gaussian tails due to leakage.

The calorimeter provides shielding from hadrons for the muon system. Eight interaction lengths (plus the muon toroid steel) are sufficient to reduce punch-through to level comparable to muons from decays and cosmic rays (see also Chapter 7).




FIG. 6-2. Events per SSC year that deposit transverse energy in a single tower greater than the threshold listed on the horizontal scale. (a) EM tower, 4 TeV $Z' \rightarrow e^+e^-$; (b) EM tower, two-jet events with $p_t^{\text{jet}} \geq 5$ TeV/c; (c) hadronic tower, two-jet events with $p_t^{\text{jet}} \geq 5$ TeV/c.

Speed and noise

The calorimeter must have sufficiently good time resolution to determine the beam crossing of electrons and jets of physics interest: $\sigma_t < 5$ ns for $E_t > 10$ GeV (e^\pm, γ) or $E_t > 50$ GeV (jets). The speed also influences the total observed excess noise (electronic plus out-of-time pileup). The total equivalent-transverse-energy noise at design luminosity should be no larger than 0.3 GeV rms in a 0.1×0.1 EM shower cone, and no larger than 3 GeV rms in a $\Delta R = 0.3$ isolation cone. An additional limit is that the contribution of noise to the missing- E_t resolution be $\sigma_{E_t} < 30$ GeV at design luminosity. Because the scintillator technology selected for the central calorimeter has very good speed and noise performance, this last requirement is largely a requirement on the forward calorimeter.

6.2.3. Requirements for the shower maximum detector

The design of the SDC detector is strongly influenced by the need to identify electrons and photons. The shower maximum detector (SMD) plays an essential role in this process by measuring the location and transverse profile of showers at shower maximum in the EM calorimeter, using narrow scintillator strips with fiber readout.




FIG. 6-3. Mass resolution for 10 TeV dijets as a function of total calorimeter depth. The jet response is derived from single-particle test beam measurements: the different curves correspond to different data sets.

SMD position resolution and strip width

The full-width-at-quarter-maximum of an EM shower is approximately 1.7 cm at the SMD. For the SMD to achieve a good measurement of the shower centroid, the shower must be distributed among several strips. The strip width must be $\lesssim 1.5$ cm. The resulting position resolution in r - ϕ or z is 2-3 mm; this is a good match to the resolution in z for tracks projected to the SMD ($\sigma_z = 2.5$ mm; Section 4.2.1).

The ability of the SMD to distinguish photons from π^0 's also depends on strip width. A study is underway to determine if this consideration sets more stringent requirements on the strip width.

SMD strip length

Cost minimization suggests that the strip lengths should be as long as allowable from physics considerations. The strip length is limited by occupancy during events of interest. The most stringent requirements are for processes in which the electron is only partially isolated, such as $t \rightarrow b \rightarrow e$. As discussed in Section 6.7, such electrons can be identified with good efficiency for strips as long as $\Delta\eta$ or $\Delta\phi \leq 0.2$.

For data collection above design luminosity, shorter strips (0.1) may be needed to reduce pileup and retain good electron identification efficiency. The shorter strips would also be a better match to the trigger towers, which are 0.1×0.1 .

SMD energy resolution

Since the SMD is not used in the precision EM energy measurement, the requirements on its energy resolution are minimal. The intrinsic resolution is set by shower fluctuations, which generate a 25% rms spread in the energy deposited in the shower maximum layer. The error on the relative energy in adjacent strips, however, is important. It must be less than 10% to avoid degradation of the position measurement. Adjacent strip pulseheight ratios can be as high as 10; this implies that crosstalk must be small enough to be understood to the level of $\sim 0.5\%$.

SMD dynamic range

The transverse energy range for electrons of interest is 10 GeV to several TeV. It is anticipated that electron identification will be most difficult at low E_t and that saturation in a strip at the very highest energies will not noticeably reduce efficiency. Thus, an adequate energy range is 10 GeV to 1 TeV. With the SMD located at shower maximum for 10 GeV electrons, the movement of shower maximum with energy produces a nonlinearity that reduces the corresponding range of mean energy deposits to a factor of 33 (see Table 6-10 below).

6.2.4. Requirements on the forward calorimeter

The forward calorimeter (FCal) extends hermetic coverage for jets to $|\eta| = 5$ (physical coverage to $|\eta| = 6$) to permit missing- E_t measurements (Section 3.4.4) and to tag high-energy forward jets, a signature of Higgs production via gauge-boson fusion (Section 3.2.3).

FCal energy resolution and transverse segmentation

The required jet E_t resolution is achievable with energy resolution $\leq 10\%$ at $E = 1$ TeV and transverse segmentation of $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$. At $\eta = 5$, the size of a hadronic shower is relatively large, and coarser segmentation is sufficient. Non-gaussian tails must be controlled to a level such that the observed missing- E_t due to the tails is not significant on the scale of physics interest, nominally 100 GeV. The radiation dose per SSC year at $|\eta| = 6$ is ~ 4000 Mrad at design luminosity. This extreme radiation dose and dose rate provides additional difficulties in achieving the required energy resolution. Such problems include potential nonuniformity due to radiation damage, space-charge or saturation effects, and pileup.

FCal speed and noise

The time resolution for jets with $E_t \geq 50$ GeV must be $\sigma_t \leq 5$ ns to allow efficient jet tagging. The speed of the calorimeter is also reflected in the pileup noise from other beam crossings. As discussed earlier, the contribution of noise to the missing- E_t resolution of the calorimeter must be $\sigma_{E_t} < 30$ GeV at design luminosity in order that the ability of the detector to make missing- E_t measurements not be compromised (Fig. 3-72). In addition, the noise in a jet cone must be consistent with the required energy resolution of 10% at 1 TeV.

6.2.5. Requirements summary

We have outlined the physics capabilities of the SDC detector in Chapter 3. This section has summarized the specific demands placed on the calorimeter. The overall calorimeter requirements are summarized in Table 6-1, while additional requirements on the SMD and forward calorimeter subsystems are summarized in Tables 6-2 and 6-3.

6.3. Calorimeter design

The previous section has described the physics requirements placed on the calorimeter. In this section we describe how these requirements are realized in the hardware. Several of the design parameters have been set using data from our R&D efforts since the submission of the Letter of Intent. In particular, two test beam efforts have provided input to the design. An electromagnetic calorimeter prototype, including shower maximum and preshower detectors, has been tested at the Fermilab MP test beam. We have also used a versatile, reconfigurable calorimeter test station (hanging file calorimeter) in which the absorber material, thickness and scintillator configuration could be easily re-arranged. Approximately 40 different configurations were explored, from a pure hadronic calorimeter using lead or iron absorbers to mock-ups of several proposed configurations for the combined SDC electromagnetic and hadron calorimeter. The details of both of these test beam programs are given in Section 6.11.

Using results from the test beam program, we have arrived at the calorimeter configuration for the central and forward detector regions previously shown in Figs. 2.6 and 2.7, respectively.

Table 6-1
Performance requirements on the SDC calorimeter.

Parameter	Requirement	Basis
η max for e^\pm ID	2.5	$H \rightarrow 4e, 2e2\mu$
EM efficiency loss in $ \eta < 2.5$	$< 5\%$	electron ID
η max for jets	5	SUSY searches
gaps in full jet coverage, $ \eta < 5$	$\leq 1\%$	Missing- E_t
EM energy resolution,		$H \rightarrow \gamma\gamma, Z' \rightarrow ee$
stochastic term	$\leq 15\%/\sqrt{E_t}$	
constant term	$\leq 1\%$	
EM transverse segmentation	0.05	$H \rightarrow 4e, H \rightarrow \gamma\gamma$
Hadronic energy resolution,		dijet mass resolution
stochastic term	$\leq 70\%/\sqrt{E_t}$	
constant term (single π^\pm)	$\leq 6\%$	
Hadronic transverse segmentation	0.10	dijet mass resolution
EM residual nonlinearity	$\leq 1\%, E_t > 10 \text{ GeV}$	$ee, \gamma\gamma$ mass resolution
Jet residual nonlinearity	$\leq 1\%/\text{TeV}, E_t > 2 \text{ TeV}$	compositeness search
Dynamic range (EM and HAC)	20 MeV–4 TeV	e ID, compositeness
EM depth	22/25 X_0	$ee, \gamma\gamma$ mass resolution
Calorimeter depth ($\eta = 0$)	$\geq 10 \lambda$	dijet mass resolution

Table 6-2
Performance requirements on the shower maximum detector.

Parameter	Requirement
Strip width	$< 1.5 \text{ cm}$
Resolution $r - \phi$ and z	$< 3 \text{ mm}$
Resolution on relative energy (strip-strip)	$< 10\%$
Strip length ($\Delta\eta$ or $\Delta\phi$)	≤ 0.2
Cross talk	$\leq 0.5\%$ after correction

Table 6-3
Performance requirements on the forward calorimeter.

Parameter	Requirement
Energy resolution	$< 10\%$ at 1 TeV
Transverse segmentation, $\eta = 3$	0.2
$\eta = 5$	0.4
Time resolution	$\sigma_t < 5 \text{ ns}, E_t > 50 \text{ GeV}$
Noise	$\sigma_{E_t} < 30 \text{ GeV}$
η coverage, jet axis	$3 \leq \eta \leq 5$
(physical)	$2.8 \leq \eta \leq 6$

Globally, the calorimeter can be considered in two pieces: 1) a central calorimeter ($|\eta| < 3.0$) in which we require lepton, photon and jet identification with excellent energy resolution and 2) a forward/backward calorimeter ($3 \leq |\eta| \leq 6$) with modest hadronic energy resolution and the ability to identify jets. The forward calorimeter is physically separated from the central calorimeter; its front face is at ± 12.5 m from the interaction point so that the transverse size of hadronic showers is more nearly matched to the required granularity of $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$. We turn first to the design of the central calorimeter followed by a discussion of the forward calorimeter.

6.3.1. The central calorimeter

The central calorimeter is divided into a central barrel and two endcap calorimeter sections as shown in Fig. 6-4. The central barrel and each endcap are segmented longitudinally into electromagnetic (EM) compartments and two different sampling hadron compartments (HAC1 and HAC2). A shower maximum detector and a so-called massless gap are integrated into the mechanical structure of the EM calorimeter. The absorber in the electromagnetic compartment is lead, and for the hadron calorimeter we have chosen iron. The barrel design consists of flat absorber plates parallel to the beam axis. The effective absorber thickness scales with polar angle (θ) as $1/\sin\theta$. It follows that the stochastic resolution term in the barrel depends only on the physically relevant variable $E_t = E \sin\theta$. In the endcap region, a smooth transition at $|\eta| = 1.4$ is made ($1/\sin\theta \sim 2$) to the barrel. Typically, absorber sampling thicknesses in the endcaps are doubled with respect to the barrel. The total depths for EM and HAC are slightly larger in the endcap due to the logarithmic increase in depth needed for the higher energy shower containment.




FIG. 6-4. Longitudinal quarter section of the central calorimeter. The physical modularity in η of barrel, endcap and removable plug is shown along with the readout segmentation in η and depth. The shower maximum detector position is indicated by the thin double line in the EM sections.

The active medium throughout the central calorimeter consists of 4 mm thick plastic scintillator tiles with wavelength-shifting fiber readout. The transverse segmentation of the calorimeter is achieved by choosing appropriate sizes of the scintillation tiles, which are arranged in projective towers pointing back toward the interaction region. The nominal EM transverse segmentation is fixed at $\Delta\eta \times \Delta\phi = 0.05 \times 0.05$, as required for good electron triggering and identification. The nominal HAC transverse segmentation is fixed at $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ for both longitudinal segments of the hadron calorimeter. The actual segmentation in ϕ is actually $2\pi/128 \approx 0.05$, and similarly, the hadron segmentation is $2\pi/64 \approx 0.1$. The endcap segmentation is the same as that for the barrel, until the point is reached where the tower size is comparable to shower size (~ 10 cm). The actual layout of the towers in the endcap is shown below in Fig. 6-42. The parameters of the central calorimeter are summarized in Table 6-4.

Table 6-4
Central calorimeter parameters.

	Barrel			Endcap			
	EM	HAC1	HAC2	EM1	EM2	HAC1	HAC2
Longitudinal readouts	1(2)*	1	1	1	1	1	1
Lateral segmentation	0.05	0.1	0.1	≥ 0.05	≥ 0.05	≥ 0.1	≥ 0.1
Absorber layers	29	28	15	6	17	20	11
Absorber material	lead	iron	iron	lead	lead	iron	iron
Absorber thickness (mm)	4.0	23.95	53.90	6.0	6.0	42.	90.
Scint. thickness (mm)	4.0	4.0	4.0	4.0	4.0	4.0	4.0
Cell thickness (mm)	10.0	30.0	60.0	12.0	12.0	48.0	96.0
Depth (not including coil)	21 X_0			6.9 X_0	18.3 X_0		
	0.85 λ	4.14 λ	4.91 λ	0.3 λ	0.8 λ	5.04 λ	5.99 λ

* Possible upgrade.

Due to differences in light output between the scintillator tiles as well as time dependent changes in the relative light output from each tile, we have incorporated a method to individually measure and mask the light output from each fiber in a tower at the photodetector ("longitudinal masking"). These longitudinal masks, when used in combination with an extensive system of radioactive sources, allow us to equalize the relative light output seen by the photodetector per unit energy deposited in each tile. Because all parts of the central calorimeter share the same readout technique, we summarize the details of the readout and masking in the next Section (6.4) on optics.

To facilitate construction, the barrel is divided into two halves with a total of 64 identical modules, as shown in Fig. 6-5. Each of the barrel modules is four electromagnetic towers wide by 28 towers long, and includes the massless gap layer, shower maximum detector, both compartments of the hadron calorimeter and readout electronics for each of these detectors. These modules are pinwheeled by 3.5 degrees to avoid radial cracks. Data from the EM module in the MP test beam indicate that with this angle of rotation the response of the calorimeter to electrons at normal incidence changes by up to $\sim 20\%$ for electrons within ± 10 mm of the interface between modules (see Fig. 6-61 below). This nonlinearity can be corrected for isolated showers with only a small degradation in resolution by using the measured relative energies between the two calorimeter modules. Both the endcap EM and hadron calorimeter are divided into 32 modules, each covering 11.25 degrees in azimuth. At the boundary between the endcap and barrel the lateral segmentation of the endcap equals the segmentation of the barrel. However, at $\eta = 3.0$ the segmentation of the endcap has grown to $\Delta\eta = 0.2$.

6.3.2. Barrel electromagnetic calorimeter

EM calorimeter depth

The barrel EM calorimeter (22 X_0 , including 1.1 X_0 from the coil, at $\eta = 0$) consists of 4 mm lead plates interleaved with 4 mm scintillator tiles and a 2 mm gap for wrapping and fiber routing. The depth has been selected such that the leakage contribution to the resolution constant term is 0.3%, after correction from the HAC1 section [3].




FIG. 6-5. Quarter section of the barrel calorimeter in ϕ . The physical modules and readout segmentation are indicated. The 3.5° pinwheel is not illustrated.

EM plate thickness

As discussed previously, detection of the decay of a light Higgs boson to two photons places the most stringent requirement on the EM calorimeter resolution. Studies [5] indicate that even for Higgs production in association with $t\bar{t}$ or W , resolution better than $\sigma_E/E = 0.15/\sqrt{E_t} \oplus 0.01$ is desirable. The predicted stochastic term for the 4 mm lead plate in the barrel calorimeter is $0.14/\sqrt{E_t}$ (including a contribution of $0.07/\sqrt{E_t}$ from photon statistics). The 6 mm plate in the endcap gives $0.17/\sqrt{E}$ for electrons at normal incidence, corresponding to $\sigma_E/E < 0.12/\sqrt{E_t}$ throughout the endcap. We have considered improvements to the resolution obtained by reducing the thicknesses of the barrel absorber plates from 4 mm to 2 mm and the endcap plates from 6 mm to 4 mm. As shown in Fig. 6-6, a 2 mm absorber thickness in the EM compartment results in a stochastic term of $9.3 \pm 0.5\%/\sqrt{E_t}$ with the same 4 mm thick scintillator [6]. The material in the solenoid coil should not prevent improved resolution due to finer sampling within the 4–2 mm range considered [7]. This option, in combination with a constant term of 1%, has been studied in Section 3.2.1. It improves the mass resolution for $H \rightarrow \gamma\gamma$ decays by approximately 25%. An improvement of 40% could be achieved by reducing the constant term to 0.5%. Such improvements to the resolution of the baseline design (with 4mm lead plates) are limited mostly by the financial resources of the collaboration.

EM longitudinal segmentation

Longitudinal EM segmentation allows electron/pion separation based on longitudinal shower shape and permits improved precision in corrections for radiation damage effects. However, studies of electron/pion separation by CDF indicate that a cut placed on the HAC1/EM ratio is correlated with the signals observed in the shower maximum detector system [8]. We might expect a similar correlation to exist with an EM2/EM1 cut, and therefore longitudinal EM segmentation is not essential for electron ID. Since the radiation damage in the barrel will be minimal, the option of two longitudinal readouts has been retained as an upgrade for the barrel, but in order to save costs, it is not included in the baseline design.

The first scintillator tile in each tower precedes the first EM absorber plate (massless gap tile) and is used to correct for energy lost by electrons in the coil (Section 6.7). This tile will be included in the optical sum with a weight of two relative to the other tiles in the tower, the increase in weight being achieved with thicker scintillator [7]. In addition, the tile will be read out separately, so that a more refined correction can be applied off-line. The weight can then be optimized for photons and electrons separately. This pre-radiator tile will also allow a statistical separation of photons and π^0 's by measuring the early shower




FIG. 6-6. Stochastic resolution term in the EM calorimeter as a function of the thickness of the lead radiator plates. The resolution includes both shower and photon counting statistics. Data for three choices of scintillator thickness are shown. The lines in the figure are to guide the eye as the relationship between the Lead thickness and resolution is not expected to be linear. The baseline design is 4 mm lead with 4 mm scintillator.

development.

A shower maximum detector (SMD) consisting of crossed narrow scintillator strips is located in the EM compartment near shower maximum. Its purpose is to measure the transverse EM shower size and position in two orthogonal coordinates. The SMD also serves, in the barrel, as a second, redundant, measurement of EM energy. This redundant feature will allow for a check against grossly mis-measured energies, due for example to HV discharges near the photocathodes ("sparks"). The physical strips are 1/8 tower wide by 1 tower long in η or ϕ . We plan initially to optically sum two strips into a logical channel 0.2 units long. We leave as a high luminosity upgrade the eventual doubling of the granularity in the SMD to a strip length of 0.1. Details of the SMD are discussed in Section 6.7.

6.3.3. EM calorimeter constant term

Critical to the success of our physics goals is the manufacture, maintenance and calibration of the EM calorimeter with a constant term in the resolution function below 1%. In Table 6-5 we list the sources of the constant term and our estimate of the contribution of each term to the resolution.

The copious production of isolated electrons from $t\bar{t}$, W and Z events should provide sufficient statistics so that the estimated error on the tower to tower calibrations is less than 0.2% [9]. For a typical electron from W or Z decay with E_t of 40 GeV, the calorimeter and momentum resolutions are about 3%. Hence, to measure the error on the mean to 0.2% requires about 300 electrons per tower. The total number of EM towers in the calorimeter is about 10,000. The relative tower to tower calibration can therefore be done with 3,000,000 electrons from W and Z production. The $W \rightarrow e\nu$ rate in the central calorimeter is 5 Hz, and 10^6 electrons can be collected in 2×10^5 s or about a week of running time at the design luminosity of $10^{33} \text{ cm}^{-2}\text{s}^{-1}$. A more detailed study can be found in Ref. 9.

Changes in the response of the photodetector/electronics chain can also be monitored by this calibration technique. Further study is needed to determine the residual constant term introduced by short-term fluctuations in photodetector/electronics response that cannot be removed by *in situ* calibration with W/Z events.

Leakage out the back of the EM calorimeter leads to another source of calibration error. While the

Table 6-5
EM calorimeter constant term budget.

Source of constant term	Contribution
Calibration tower to tower	0.2%
Leakage	0.3%
Transverse uniformity	0.5%
Tile-to-tile variations incl. thickness variations and longitudinal masking	0.5%
Absorber thickness variations	0.2%
Radiation damage	0.5%
Total (added in quadrature)	< 1.0%

energy deposited in the HAC1 calorimeter can be used to correct the energy measured in the EM section, there remains a residual energy dependent contribution of 0.3% [3].

In Section 6.4 (optical system) we demonstrate that transverse nonuniformity of light collection across a tile can be maintained to better than 2%. This variation, assumed to be completely coherent, is shown to correspond to less than 0.5% in the constant term. Random tile-to-tile light output variations due to variations in the tile thickness, optical path or similar effects, masked to $\leq 2\%$, lead to a contribution to the constant term of 0.5% or less.

The allowed variations in absorber thickness are $\pm 2\%$ rms. These random errors lead to a constant term of $< 0.2\%$ in the electromagnetic calorimeter. This standard can be met with commonly available manufacturing techniques.

Finally, for the barrel calorimeter the maximum radiation dose is about 60 krad in 10 years of running at design luminosity (see Fig. 6-7a). In Section 6.5 (radiation damage), we estimate that this radiation dose contributes less than 0.5% to the constant term for the barrel (see also the discussion of radiation dose and damage for the endcap, below). Furthermore, radioactive source tubes integrated into the mechanical structure of the calorimeter allow us to measure the light output from each tile. Hence, for longer time scales we can correct for the radiation damage in the barrel using the source tube data to calculate appropriate longitudinal masks that equalize the light output for each tile. We discuss radiation damage in the endcap electromagnetic and hadron calorimeter below.

6.3.4. Endcap electromagnetic calorimeter

The endcap EM calorimeter is constructed of 6 mm lead plates, 4 mm scintillator tiles, with a 2 mm gap as in the barrel, and is $25 X_0$ thick. Its design is similar to that of the barrel except for the modularity of the mechanical design. However, two issues distinguish the endcap EM calorimeter from the barrel: radiation damage and activation.

In the endcap EM sections the radiation field is intense. In Fig. 6-7 we show the expected radiation dose from minimum bias events, including both electromagnetic and hadronic contributions [10]. The contours in Fig. 6-7b were estimated using electron and pion longitudinal shower profiles measured at 5 GeV (corresponding to $p_t \approx 0.7$ GeV/c) in the hanging file calorimeter. At any given η , the maximum radiation dose is four times higher in the EM compartment than in the hadron compartment.

The implementation of two EM longitudinal compartments in the endcap allows one to make radiation-damage corrections for isolated electrons and photons from event-by-event measurements of the shower initiation point. For example, a radiation dose of 1 Mrad corresponds in our radiation damage tests to a




FIG. 6-7. (a) Maximum ionizing dose in the SDC calorimeter for one year of running at design luminosity and (in parentheses) for 10 years at $10 \times$ design luminosity [1]. (b) Equal radiation dose contours for the SDC endcap calorimeter after ten years of running at design luminosity. We have indicated in (b) the separation between the EM and hadronic compartments and overlaid several η boundaries.

light loss of about 30% at electromagnetic shower maximum. Knowledge of the interaction point of the electron or photon (derived from the EM2/EM1 ratio) and of the degree of signal loss at shower maximum (revealed by the calibration system and by the average response of the SMD) allows much of the resolution loss to be recovered as shown in Fig. 6-8 [11].

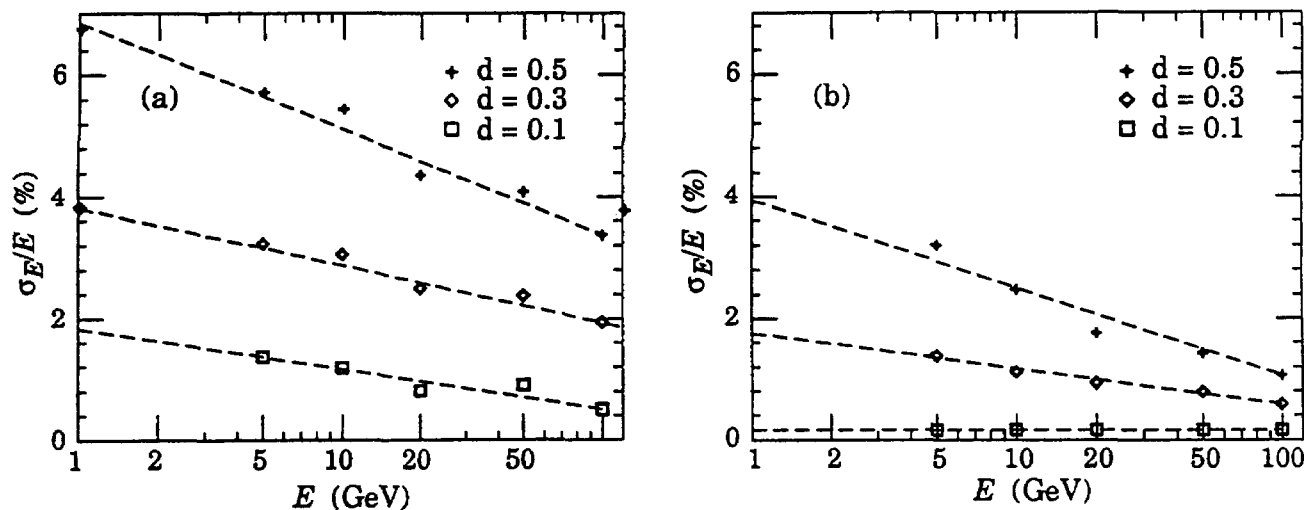


FIG. 6-8. Radiation damage contribution to the electron resolution as a function of energy for several radiation damage parameters (d is the fraction of light lost at shower maximum) (a) before correction (b) after correction with EM2/EM1 ratio.

The absorber material of the endcap will become activated. In Fig. 6-9 we show contour plots of the activation of the endcap after 30 days of running at design luminosity followed by one day of cool-down.

The activation is highest near the interface between the EM and HAC1 compartments, reaching a level of about 20 mrem/hr on contact. These values are not daunting by accelerator standards, although temporary shielding will have to be provided to work in the immediate vicinity.




FIG. 6-9. Equal activation contours in units of mrem/hr on contact for the SDC endcap calorimeter after 30 days of running at design luminosity followed by one day of cool-down. The figure was generated assuming the front face of the EM calorimeter begins at $z = 4.2$ m from the interaction region. Longer periods of exposure at the same average luminosity would increase the values by factors of 2 to 3. Extending the cool-down period from one day to one year would reduce the values by a factor of 3. We have indicated separation between the EM and hadronic compartments and overlaid several η boundaries.

The implementation of two longitudinal EM segments implies that the endcap electromagnetic calorimeter will require refurbishment only on the time scale of a few years in order to maintain a resolution constant term within SDC specifications. Our baseline design for the endcap EM calorimeter allows for the replacement of the scintillator by dismounting the EM calorimeter in wedge shaped modules each corresponding to $2\pi/32$ of azimuth. Similar considerations make it prudent to plan to repair and refurbish the endplug region of the endcap HAC compartment, $\sim 2 < |\eta| < 3$. We are studying other arrangements that would also allow refurbishment.

6.3.5. Barrel and endcap hadron calorimeter

Our design provides for two compartments in the hadron calorimeter with different absorber thicknesses. The hadronic section closest to the EM calorimeter (HAC1) has finer sampling thickness to accurately measure jet energies. These thinner iron plates also aid in measuring electromagnetic energy leaking out of the rear of the EM calorimeter to improve the energy measurement of high energy, isolated electrons. The second hadronic calorimeter compartment (HAC2), with twice the sampling thickness of the HAC1 section, is used to catch the energy in the tails of high energy showers and to estimate leakage out the back of the calorimeter. The division of the hadron calorimeter into two compartments with different sampling fractions allows us to meet our resolution and depth requirements at lower cost than would be possible with a calorimeter having uniform sampling.

Hadronic calorimeter parameters

The main hadron calorimeter parameters are listed in Table 6-4.

Hadronic depth and longitudinal segmentation

As discussed in the calorimeter physics section, we have chosen the total calorimeter depth to be 10 interaction lengths at $\eta = 0$ to provide sufficient containment for high energy jets. We have verified that division of the hadronic calorimeter into two compartments does not significantly degrade the resolution using data from the hanging file test beam calorimeter. In Fig. 6-10 we show the resolution for 227 GeV pions as measured by the hanging file test beam calorimeter as a function of the combined thickness of the EM and HAC1 sections. The total thickness is held at 10 interaction lengths. The resolution increases from 5.3% for all 1-inch iron absorber to 5.7% for a 4.9 interaction length HAC2 constructed with 2-inch iron plates.




FIG. 6-10. Resolution (σ_E/E) for 227 GeV pions as a function of the EM+HAC1 depth, with various HAC2 absorber thicknesses. Total calorimeter thickness is 10 λ . These data are from the hanging file calorimeter measurement.

Hadronic absorber material

Iron has been selected as the absorber for the hadron calorimeter after an evaluation of both iron and lead on the basis of performance for single pions and jets, impact on other systems, engineering feasibility and cost. The two specific absorber configurations evaluated are outlined in Table 6-6. They differ in the absorber used in the HAC1 compartment and in the thickness of HAC1; the EM compartments are identical, as is the unit cell used in HAC2.

Lead versus iron performance comparison

The linearity and resolution of the specific configurations under consideration have been evaluated using CALOR89 [12], and similar configurations have been studied with the hanging file test calorimeter at Fermilab. The results are summarized in Table 6-7 and Fig. 6-11.

Table 6-6a
Simulation model A (iron HAC1)*.

	EM	Barrel HAC1	HAC2	EM	Endcap HAC1	HAC2
Longitudinal readouts	1	1	1	2	1	1
Lateral segmentation	0.05	0.05	0.1	≥ 0.05	≥ 0.05	≥ 0.1
Absorber layers	36	32	12	18	21	7
Absorber material	lead	iron	iron	lead	iron	iron
Absorber thickness (mm)	3.2	25.4	51	6.3	51	102
Scint. thickness (mm)	4	2.5	2.5	4	2.5	2.5
Depth	$21 X_0$	4.8λ	3.6λ	$21 X_0$	6.4λ	4.3λ

* The configuration described in this Design Report is similar to model A, but differs in several parameters.

Table 6-6b
Simulation model B (lead HAC1).

	EM	Barrel HAC1	HAC2	EM	Endcap HAC1	HAC2
Longitudinal readouts	1	1	1	2	1	1
Lateral segmentation	0.05	0.1	0.1	≥ 0.05	≥ 0.1	≥ 0.1
Absorber layers	36	24	18	18	24	10
Absorber material	lead	lead	iron	lead	lead	iron
Absorber thickness (mm)	3.2	21	51	6.3	31.5	102
Scint. thickness (mm)	4	2.5	2.5	4	2.5	2.5
Depth	$21 X_0$	3.0λ	5.5λ	$21 X_0$	4.4λ	6.1λ

Table 6-7

Linearity and resolution coefficients $a/\sqrt{E} \oplus b$ for single pions predicted by CALOR89 and measured by the hanging file test calorimeter. The integration time (τ) is also given.

Case	τ (ns)	e/π at 30* GeV	$a\sqrt{\text{GeV}}$	b
CALOR89				
Model A barrel (iron)	16	1.066 ± 0.017	0.60 ± 0.03	0.058 ± 0.005
(same)	96	1.055 ± 0.014	0.57	0.048
Model B barrel (lead)	16	1.023 ± 0.013	0.57	0.042
(same)	96	1.00 ± 0.010	0.52	0.038
Model A endcap	16	1.069 ± 0.013	0.73	0.055
(same)	96	1.045 ± 0.013	0.67	0.042
Model B endcap	16	1.020 ± 0.013	0.74	0.052
(same)	96	0.963 ± 0.013	0.63	0.071
HANGING FILE preliminary				
Iron hadron section	96	1.020 ± 0.011	0.51	0.037
Lead hadron section	96	0.991 ± 0.011	0.49	0.049

* The calibration is such that electrons and pions of 300 GeV are correctly measured.

Both the hanging file data and the CALOR89 calculations indicate that the lead HAC1 compartment does not result in substantially better resolution. The large constant term is largely due to the difference in sampling frequency between the EM and hadron compartments. The mismatch in compensation properties of the EM and HAC1 sections is illustrated in Fig. 6-12, which shows the measured energy as a function of the energy in the EM compartment.

ISAJET two-jet events were used to evaluate the jet response. The single particle linearity and resolution (σ_E/E) from CALOR89 were tabulated as a function of energy and used with a parametrized shower generator to evaluate the jet response. The results are summarized for the barrel calorimeter in Table 6-8. As expected, the constant term is reduced by the averaging effect of multiple particles. Neither the iron nor the lead HAC1 give a linear response, but the correction is somewhat smaller for lead. The jet resolution in the endcap is less dependent on energy. For $2.0 < |\eta| < 2.5$ and $E_t \geq 200$ GeV, $\sigma_E/E \approx 0.030$ for model A (iron HAC1) and 0.022 for model B (lead HAC1). Note that the lesser energy dependence in the endcap is for fixed E_t ; E is large and the constant term dominates.

Table 6-8

Jet response (barrel), $a/\sqrt{E} \oplus b$, from a CALOR89 simulation.

Case	τ	nonlinearity*	0.1 TeV	at 1 TeV	at 10 TeV	$a\sqrt{\text{GeV}}$	b
Model A	16 ns	0.918 ± 0.002		0.962	0.998	0.633 ± 0.002	0.019
Model A	96 ns	0.922		0.964	0.998	0.56	0.016
Model B	16 ns	0.967		0.982	0.998	0.52	0.014
Model B	96 ns	0.997		1.00	1.00	0.44	0.010

* Nonlinearity defined as $E_t(\text{measured})/E_t(\text{true})$.

The response in Table 6-8 does not include the effects of fluctuations in and out of the jet cone. These fluctuations are important for mass resolution, but are less so for compositeness studies, where a comparison is made to the predicted jet cross section including the fluctuations.




FIG. 6-11. (a) e/π vs energy as predicted by CALOR89 for 16 ns integration time for iron HAC1 (solid) and lead HAC1 (dashed). (b) e/π measured by the hanging file calorimeter (96 ns integration).




FIG. 6-12. Average energy measured for 300 GeV pions vs energy deposited in the EM compartment. CALOR89 calculation is for model B; model A shows similar behavior.

Other considerations of absorber material choice

As an absorber material, iron has two substantial advantages over lead—its magnetic properties and its lower cost. The iron HAC1 terminates the magnetic field more effectively than the lead. This produces a more uniform field in the tracking volume (1.3 T vs 1.0 T at $z = 4$ m) and substantially reduces the compressive load on the solenoid (from 1700 tons to 1150 tons).

We have considered a variation of model A—the “reentrant endcap”—which would place the iron of HAC1 very close to the current sheet of the solenoid, producing a much more uniform field and reducing the compressive load on the solenoid to 310 tons. However, the reentrant endcap has a break in the EM calorimeter that is not present in either of the other two configurations studied. Such a discontinuity would lead to additional loss of fiducial volume, and this choice was not adopted.

Summary of hadronic calorimeter absorber

Simulations indicate that model B (lead HAC1) should have slightly better resolution and linearity than model A (iron HAC1). Preliminary test beam data indicate that an iron HAC1 in a configuration close to, but not identical to model A, has practically the same resolution as a lead HAC1. In either case, the differences have not been shown to have significant physics impact.

The iron HAC1 section significantly reduces the force on the solenoid and has a lower cost. It is also better in terms of the impact on the EM calorimeter [3] and the tracking system (lower neutron flux) and provides a more uniform magnetic field. The SDC has, therefore, selected iron as the absorber for the HAC1 compartment for the central calorimeter.

6.3.6. Forward calorimeter (FCal)

As calorimeter coverage is continued to the region of η beyond 3.0, (to angles smaller than 100 milliradians), the combined influence of several factors requires the front face of the calorimeter to be further away from the interaction point. In the endcap calorimeter, whose front face is at $z = 4.5$ meters, hadron shower size has already begun to dictate the tower granularity in η and ϕ near $\eta = 3$. The forward calorimeter must maintain a smooth and hermetic connection to the endcap calorimeter, but it is also highly desirable to provide a segmentation $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$. Over a significant part of the forward calorimeter η range, this implies that the front face must be at $z > 12$ m in order to keep tower sizes larger than showers.

Moving the front face away from the interaction point helps with several other problems: it reduces the tower occupancy rate and hence the tendency towards pileup; it reduces the load on the tracking detectors caused by albedo from the calorimeter front face; and it reduces radiation damage and radioactivation of calorimeter elements.

Forward calorimeter geometry

The transition from the endcap to the forward calorimeter is chosen to occur at $\eta = 3.0$. At this boundary, a conical hole is cut in the endcap calorimeter, and this opening in the acceptance is covered by the forward calorimeter, whose front face is at $z = 12.5$ m. The FCal geometry is shown in Fig. 6-13. The forward calorimeter is made oversize radially by one interaction length to improve the smoothness of the endcap/forward calorimeter transition. This radial extension provides lateral containment for showers initiated by particles striking the forward calorimeter front face at $\eta = 3.0$, and it helps catch spray from showers which have already begun close to the endcap calorimeter's conical aperture. The transverse segmentation is shown in Fig. 6-14.

The forward calorimeter high rate environment

The severe rate environment for the forward calorimeter places special constraints on the design of this subsystem. The forward calorimeter is particularly susceptible to radiation damage. Although the particle multiplicity for a given $\Delta\eta \times \Delta\phi$ area is independent of η , particles emitted at high η are boosted to very high energies. They tend to produce denser, more extensive showers, and collectively they produce more radiation damage. Moreover, a $\Delta\eta \times \Delta\phi$ area (0.1×0.1 for example) shrinks as θ^2 in this region. The SDC "lifetime radiation dose" (taken as 10 years at $10^{34} \text{ cm}^{-2}\text{s}^{-1}$) at high η rises well into the Gigrad region, as shown in Fig. 6-13. This value is so high that radioactivation of forward calorimeter components must be considered from an ultimate environmental/disposal point of view, and the radioactivation of both forward calorimeter and SSC machine components must be considered from the point of view of safety of personnel engaged in forward calorimeter servicing activities.

Much of the physics requiring the highest η coverage has relatively high rate and can be accomplished during design luminosity running. The part of the forward calorimeter occupying the smallest angle region ($\sim \eta > 4.5$) could if necessary be removed as the luminosity grows to $10^{34} \text{ cm}^{-2}\text{s}^{-1}$; the search for a light gluino, for example, might have been completed. Such a strategy would, however, abandon a significant part of the capability for jet tagging in a hunt for a heavy Higgs produced by WW/ZZ fusion.




FIG. 6-13. Forward calorimeter geometry. Radiation doses are for 10 years at $\mathcal{L} = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. Note the required clear aperture goes beyond $\eta = 3.0$.




FIG. 6-14. Tower layout in the forward region. Note the "seamless" boundary between the central calorimeter and the forward calorimeter at $\eta = 3.0$. At the endcap front face this boundary is 0.45 m from the beam line, and at the forward calorimeter front face it is 1.25 m from the beam line.

6.4. Tile/fiber optical system design

Scintillation tiles read out by wavelength shifting fibers are used for sampling the energy deposition in the central calorimeter. The radiator plates are interspersed with scintillator tiles, whose transverse dimensions determine the transverse segmentation of the calorimeter. The design uses tiles whose front face dimensions (in the EM sections at $\eta = 0$) are approximately $11.0 \times 11.0 \text{ cm}^2$ to yield a segmentation of 0.05×0.05 in $\Delta\eta \times \Delta\phi$ space. The scintillation light from energy deposited in the scintillator plates is collected by wavelength shifting fibers embedded in a groove in the tile and is thereby transmitted to the photon transducer, a green-sensitive photomultiplier tube (PMT). To maximize light transmission to the photon detector, the waveshifting fiber is fused to a clear plastic optical fiber close to the point at which it exits the tile. The use of fibers rather than bars or plates is dictated by the desire to maximize the uniformity (no "hot spots") and hermeticity of the calorimetry. Wrapping of the tiles with aluminized mylar, painting

of the edges of tiles with reflective white paint, and masking near the photocathode are used to maximize light output and achieve the required transverse and longitudinal uniformity.

Much work has been carried out to determine the optimum fiber path within a tile. We propose to use the "sigma" path in which one end of the fiber is mirrored and the other end is read out (Fig. 6-15), instead of reading out both ends of the fiber. This approach has recently been implemented in a test calorimeter, with considerable success. Preliminary data indicate that an acceptable light yield and transverse uniformity can be achieved. The resulting decrease in the readout fiber bundle eases mechanical logistics and improves the hermiticity of the calorimeter.




FIG. 6-15. Sketch of sigma tile fiber route.

One limitation of the tile/fiber technique is the minimum radius through which plastic fibers can be bent. This is approximately 2.5 inch for 1 mm diameter fibers. We anticipate no difficulty in instrumenting the tile sizes required to achieve the proposed barrel transverse segmentation.

6.4.1. Light yield

We address only EM calorimeter light yield requirements since the required light yield for a hadron calorimeter is much less, given the required resolutions discussed in Section 6.2.

A stochastic EM resolution term $\leq 15\%/\sqrt{E_t}$ is required. In the sampling scintillator plate calorimeter presented here, this resolution depends on: the intrinsic light output of the scintillator and waveshifter, the thickness of radiator and scintillator plates, the diameter of the wavelength shifting fiber and its route within the tile, and the quantum efficiency of the photon transducer. A light yield of greater than 200 photoelectrons/GeV or 2 photoelectrons/minimum-ionizing particle (MIP)/tile is required.

We have measured the relative light yields of a number of common scintillator plastics. Normalizing to SCSN-81, our current choice, the observed light yields are as follows: SCSN-81 = 1.0; SCSN-38 = 1.22; BC-400 = 1.45; BC-404 = 1.84; BC-408 = 1.70. Although the Kuraray plastics SCSN38 and 81 show less light yield than the pure PVT Bicron plastics BC404 and 408, they are more radiation resistant. For this reason, we are proposing the use of SCSN-81, which is the most radiation hard. However, in our R&D program, we are actively evaluating alternative plastics and fluor systems to determine if systems with greater radiation hardness and higher light output can be fabricated.

The thickness of radiator and scintillator plates best matched to the required resolution was determined by analysis of data taken with a tile-fiber EM calorimeter test module [14]. In this test, the energy dependence of the resolution for two values of light output was used to determine the intrinsic light yield and a scaling parameterization based on this and CDF data were used to compute the expected resolution as a function of radiator and scintillator plate thickness (Fig. 6-16). A calorimeter using 4 mm radiator and 4 mm scintillator plate thickness is matched to the upper limit of acceptable stochastic resolution.




FIG. 6-16. Scaling of the stochastic term in the EM shower counter as a function of radiator thickness, scintillator thickness and assumed light yield. Data are shown from the ANL/WSTC Test Module and the CDF Central Electromagnetic Shower Counters. The curves are predictions for 2.5 mm scintillator (solid), 2.5 mm with $\times 0.5$ light yield (dotted), 4 mm scintillator (dashed), and 4 mm $\times 0.5$ light yield (dot dash).

Better resolution is possible with thinner radiator, but this option has not been exercised due to cost considerations.

The maximum fiber diameter is determined by the minimum required bend radius of 2.5 cm. Thus we are developing a design based on 1 mm diameter fiber. We have studied the light yield as a function of the fiber layout on the tile. Figure 6-17 shows the light yield in photoelectrons for a selection of layouts varying from a single-ended sigma tile to a spiral layout in which 3 rings of fiber are embedded in the tile [15]. The circles in this figure correspond to tiles spliced to 4 meters of clear fiber, the worst case for the SDC calorimeter. The sigma tile has 4 p.e./MIP in this configuration for 4 mm thick BC 408 scintillator. Similar results were obtained in an independent series of measurements described in Ref. 17.

The emission spectrum of a typical green waveshifter, Y7, is shown in Fig. 6-18. Also shown are the quantum efficiency curves for a standard bialkali and an extended green photocathode. We see that the latter has a better quantum efficiency for most of the emission spectrum and would generate more photoelectrons. The CDF Plug Upgrade Group found that a prototype tile with BCF-91 fiber coupled to 4 meters of clear fiber yielded 5 p.e./MIP using a bialkali phototube and 9 p.e./MIP with the Phillips XP2081 extended green photocathode. Hence, a 4.0 mm thick SCSN-81 sigma tile coupled to a XP2081 phototube after 4 meters of clear fiber is expected to yield ~ 4.0 p.e./MIP. The Hamamatsu R580-17, an alternative extended-green sensitive phototube, is also being evaluated.

6.4.2. Longitudinal light yield variation

We have done Monte Carlo studies to determine the maximum tile-to-tile variation in light yield consistent with the required constant term in the calorimeter resolution.

The relationship between random plate variations and the hadronic constant term was determined using Fermilab Lab E test beam data [18]. CCFR pion data of various energies were studied. The CCFR calorimeter used 10 cm thick iron absorber plates interleaved with scintillator plates, with each scintillator plate being read out by a separate phototube. After removing the fluctuations the plane-by-plane data were reanalysed including a hypothesized $60\%/\sqrt{E}$ sampling term expected for 2.5 cm iron sampling and a longitudinal tile-to-tile variation in response characterised by a gaussian distribution. The constant term as a function of rms light yield variations for several energies is shown in Fig. 6-19.




FIG. 6-17. Light yields for several fiber layouts: 1) sigma tile; 2) U tile; 3) spiral tile with two loops; 4) spiral tile with three loops. The open circles are light yield after 4 meters of clear optical fiber. The light was collected by a bialkali phototube.




FIG. 6-18. Emission spectrum for Y7 (solid curve). Quantum efficiency spectrum for bialkali photocathode (dots) and Phillips 2081 extended green photocathode (dashes) are also shown.

The relationship in Fig. 6-19 is linear; a 2% constant term corresponds to 5% rms random scatter in longitudinal response. Therefore, the hadron calorimeter optical system being developed allows a 5% random longitudinal variation in light yield in order to maintain a contribution to the constant term of less than 2%.




FIG. 6-19. Constant term in resolution (b) vs tile-to-tile light yield variation for hadron showers. Test beam data from the "Lab E" spectrometer were used at beam energies of 0.1, 0.2, and 0.45 TeV.

For electrons, the maximum allowed constant term from all sources, 1%, is a more stringent requirement. The EM calorimeter structure was simulated using an EGS Monte Carlo, and the induced constant term has been extracted for various random variations in absorber and scintillator plate thicknesses [19]. The results are shown in Fig. 6-20. The conclusion is that the scintillator plate response must be controlled to an rms of 2% if the induced constant term is to be held within the SDC calorimeter specification, taken

here to be a 1/2% induced term. This study also included the effects of tower-to-tower normalization assuming some fraction of the modules are calibrated in test beams but the majority normalized using sources. These effects become irrelevant once *in situ* data calibration becomes possible; the 2% and 5% requirements are thus conservative.




FIG. 6-20. Constant term in resolution (b) vs tile-to-tile light yield variation for EGS simulated electromagnetic showers for 10 GeV electrons (open diamonds) and 50 GeV (solid diamonds).

6.4.3. Transverse light yield variation

Transverse variation of tile response will also affect the constant term in the calorimeter resolution. Because electron showers are much smaller than hadron showers, they probe more critically the required transverse uniformity. An EGS Monte Carlo incorporating measured tile nonuniformity maps was used to study the required uniformity. In the model, 100% coherent nonuniformity was assumed (the worst case scenario). Figure 6-21 shows the measured transverse response from a CDF prototype tile [20]. Before masking, the rms variation is about 7%; this is reduced to about 2% by a simple black/white paper mask. Figure 6-22 shows the energy resolution vs $1/\sqrt{E}$ for an EGS study using these measured transverse responses. The uncorrected response yields a 1.2% constant term. This is reduced to 0.4% by use of the response-flattening masks. We conclude that an rms variation across the tile face of 2% or less is acceptable.

6.4.4. Optical system fabrication studies

We now turn to the mechanical issues associated with realising uniformity in the production of a large calorimeter.

Splicing

There has been an extensive R&D program on splicing polystyrene fibers [21]. Two prototype "clam shell" splicers, operated manually, were used in the construction of the CDF endplug prototype and the SDC EM test module. Figure 6-23 shows results on the transmission through the splice. We see that the rms variation is about 4%, which can be readily masked. Our continuing R&D program includes construction of an automated splicer, which should achieve better performance. However, if required, splicing variations can be removed at the last stage of assembly through longitudinal response measurement and masking.




FIG. 6-21. Transverse scan (using radioactive source) for a "U" tile before and after masking the tile surface with black/white paper to flatten the response. The scan crosses the fibers of the U.




FIG. 6-22. Energy resolution for electrons vs $1/\sqrt{E}$ from an EGS simulation using measured tile transverse nonuniformities as described in the text. The induced constant term reduces from 1.2% in the uncorrected case (open points) to 0.4% for tiles whose response has been flattened using a simple mask (solid points.)

Transverse uniformity & generic masking

Fiber placement is crucial to transverse uniformity of the tile. In general, we have found that a closed loop of fiber is better than an open "U." The transverse response for a sigma tile wrapped with aluminized mylar is shown in Fig. 6-24. The data were taken using a collimated ^{106}Ru source. The pulse height distribution for β 's penetrating the tile and two trigger counters was formed as a function of position on the tile [17]. The small beam spot of the source (2×5 mm) allowed us to probe the transverse uniformity at a scale smaller than an EM shower. The measured nonuniformity of 2% rms demonstrates that, for this choice of fiber route, no masking is required to meet the design specification. We have considered other layouts than the sigma option, as there may be mechanical or other issues that would prevent use of sigma tiles throughout the calorimeter. We have therefore investigated a general procedure for flattening the responses of tiles with arbitrary initial nonuniformity using the common technique of masking the wrapper of tiles and having black areas on the wrapper where we want to reduce the light yield. We have found




FIG. 6-23. Splice transmission variation for a production run of splices used in the CDF upgrade calorimeter prototype.




FIG. 6-24. Transverse response map for a sigma tile.




FIG. 6-25. Tile-to-tile light yield variation for tiles in the SDC EM test calorimeter. The rms variation is 8.1%.

that for the geometry of tiles tested (2.5 mm thick SCSN 81) we can flatten the tile response to a few % if the maximum initial nonuniformity in response is less than 20%.

Tile light yield variations, longitudinal uniformity and masking

As shown above, tile-to-tile light yield variations must be controlled to 2%. This is a stringent specification on tile-fiber fabrication. Figure 6-25 shows the light yield fluctuations for a typical set of tiles for the SDC EM test calorimeter. An rms variation in light yield of 8.1% is observed. Only limited improvements in quality via automation of tile assembly may be expected. Therefore, we include in the optical system design a system that allows us to correct tile-to-tile light yield variations after the calorimeter has been constructed.

Tiles are installed into the absorber of the calorimeter and the fibers are dressed to the appropriate phototube cookies. The cookie has a hole for the fiber from each tile of the tower. An extensive radioactive source calibration system is built into the calorimeter so that each tile can be exposed to the same source

sequentially. The current generated by the phototube when the tile is exposed to the source is a measure of the response of the tile and the optical transport to the phototube. A perfect tower would have identical responses for each tile in depth (except for solid angle effects in capturing energy from the source). Following the measurement of response, a computer-generated neutral density filter is placed between the cookie and the phototube to correct for any observed depth variation of the tiles.

Figure 6-26 shows a computer generated mask to be converted into a neutral density filter. Figure 6-27a is a longitudinal source scan of a typical tower in the electromagnetic calorimeter of the CDF prototype. On the horizontal scale is the tile number ranging from 1 (front of the tower) to 23 (back of the tower). The vertical scale is the relative light yield of the tiles. We see that the 10.5% rms variation (before masking) has been reduced to 2.1% (after masking). We note that 2% longitudinal variation is within the SDC specification.




FIG. 6-26. Computer output used to form a neutral density filter for longitudinal masking.




FIG. 6-27. The longitudinal response of tiles in a tower. The horizontal scale is tile number ranging from 1 (front) to 23 (back). The unmasked RMS variation is 10.5%. Following longitudinal masking the RMS variation reduces to 2.1%.

Mirroring

Fiber mirroring techniques were developed to be employed on the sigma tile. The mirrored end replaces a second fiber readout, and, if high reflectivity can be obtained, yields comparable light and uniformity to that obtained for a design in which both fiber ends are coupled to a photon detector.

Fibers were cut on a diamond fly cutter specifically developed to cut fiber ends perpendicular to the fiber with a smooth finish. The fiber end to be mirrored was washed in 50% ethyl alcohol and then aluminum was vacuum-evaporated onto it. We studied several thicknesses of deposit: 200, 400, 800 Å. The reflectivity was found to be independent of the coating thickness. A reflectivity of 0.84 ± 0.03 was obtained for an 800 Å coating of aluminum, which is comparable to results obtained elsewhere [22]. We have subsequently explored the radiation hardness of the mirroring finish. Figure 6-28 shows the mirror reflectivity measured as above for sets of fibers exposed to various radiation doses with a ^{60}Co source. We see no radiation-induced degradation for doses up to 3 Mrad.




FIG. 6-28. Fiber mirror reflectivity vs radiation dose for mirrored fibers exposed to a ^{60}Co source. Doses of 0, 0.1, 1, and 3 Mrad are shown.

Fiber aging/annealing

A source of possible concern is fiber aging such that the optical properties of the fibers change as a function of time. If such an instability existed, calibration would be more difficult. We have initiated a program of fiber aging studies to check if there is such an effect and, if so, how to minimize it.

Tiles from an SDC tile prototype EM calorimeter [15], constructed in the summer of 1990, were periodically retested to see if there was any change in light yield. In addition, the fiber transmission through these tiles was also retested. We have not detected aging effects so far, after about 14 months. The fibers in these tiles were not pre-bent or treated in any special manner before installation.

We have also done tests on fibers not embedded into tiles aimed at amplifying any small aging effects that may be unobservable in the tiles themselves, in particular those associated with fiber bending and handling. We have built coils of fiber (4.5 and 9 times 360° rotation) and studied the change in transmission through the fiber as a function of time. Various initial treatments (annealing) of the fiber coils were studied to determine if they provided better long term aging. The techniques studied included prebending the fiber in hot air, in hot water baths of varying temperatures and for several cool-down patterns. The loops were measured over a period of about 40 days. Radii of curvature of 0.5 and 1 inch were studied. We find that no annealing (simply wrapping the fibers into the coil) yields the best aging properties. For these coils, after an initial light loss occurring within several minutes of forming the coil, no additional aging is observed.

Tile cutting

The baseline procedure for cutting the tiles and fiber grooves is high-speed CNC machining. The focus has been on an upgraded version of the machine currently used to cut tiles for the CDF endplug upgrade. The machine is equipped with a CNC control that can be downloaded directly from a CAD database, greatly simplifying the task of producing the large variety of tile sizes and making it very cost effective. It offers a conservative approach and produces an acceptable edge and groove finish. The current groove geometry is U-shaped, but other shapes, such as a ball-groove, which would hold the fiber in place, are being considered.

Tests have been carried out on samples of polystyrene to optimize cutting speed and surface quality. Results show that the tiles can best be cut out by using a high speed steel endmill at a rate of 175–200 in/min. and 17,000 rpm. The grooves are easily cut using a high speed steel, single flute, ball endmill at a rate of 75 in/min and 18,000 rpm.

Edge painting

The current design baseline calls for painting the edges of all tiles with white reflective paint in order to enhance the light yield uniformity. The focus has been on finding techniques which minimize handling and therefore cost. The requirement for painting the edges as well as the specific techniques and materials to use are under active study. At present, the baseline design envisages a robotic painting arm used in a way commonly found in industry to paint stacks of approximately 100 tiles in one operation.

Waveshifter fiber insertion

The U-groove geometry requires that some means be provided to hold the fiber in the groove. The baseline proposal is modeled on the CDF approach where the fiber is held in several places with clear glue. Efforts are being made to find optically clear, radiation hard glues which will not adversely affect the fiber cladding. Several alternative methods of fastening the fiber have been proposed and are being studied. These include modification of the groove shape to a keyhole shape in which the fiber is mechanically located.

6.4.5. Photon detector and associated electronics

A primary design goal for the transducers and amplifiers is stability. The time scale for stability is on the order of days. This performance appears to be readily achievable [16].

Other requirements are dynamic range (20 MeV to 4 TeV is sufficient), small non-linearity after correction ($< 0.25\%$), and precision ($< 0.25\%$). The main problem is that the large dynamic range, coupled to the maximum output current, implies that the PMT must be operated at a rather low gain, a mode of operation somewhat different from standard procedure. As subsidiary conditions, the package size must be fairly small, and an enhanced green response is most desirable. A draft set of specifications is listed in Table 6-9.

Several candidate PMTs exist that meet these specifications. The technology is not unique nor does SDC push its limits. There is an existence proof [16] of a similar PMT base system.

6.5. Radiation hardness of tile/fiber calorimetry

The tile/fiber calorimetric detection technique emphasizes speed of response and hermeticity. Stability and uniformity of response must be monitored, and relative calibration must be maintained if the SDC goals for calorimetry are to be met. The radiation field of the SSC, shown in Fig. 6-7, is of unprecedented intensity for collider experiments. Radiation damage to the plastic tiles and fibers is potentially a serious problem, as the calorimeter signals will decrease with exposure.

This varying detector response has several implications for the design. As discussed below in Section 6.8, one is the necessity to incorporate a radioactive source tube monitoring and calibration system that will measure the damage profile by recording the response of each tile to a radioactive source over the duration of the experiment. The source measurements will also be an important part of the quality control procedures during manufacture. In particular, they will provide the final gain measurement to determine neutral density filter masking for longitudinal uniformity. A source calibration scheme has been used extensively in beam tests undertaken to study radiation damage [23]. It provides a 1% relative measurement of the light output of all the tiles in a tower. Another design choice, motivated by radiation damage concerns, is the longitudinal segmentation into two EM compartments in the endcap. The ratio of the signals in the two EM sections determines the longitudinal location of the origin of an isolated EM shower. This information allows the shower response to be corrected according to the damage profile measured with the source tube system. The atmosphere surrounding the tiles during damage and annealing is also of importance [24]. For this reason, the design incorporates a gas-tight structure which allows the atmosphere surrounding the plastic array to be controlled during irradiation and during recovery. Finally, provision is made for replacing damaged plastic in regions of the endcap.

Table 6-9
Draft PMT specification for SDC calorimetry.

Parameter	Requirement
Mechanical	
Diameter	19 mm (3/4") to 40 mm (1.25")
Length	~75 mm
End window	Flat face
Leads	Bare pins, solder leads, or socket type
Photocathode	
Useful Area	1 cm ² minimum good area
Quantum efficiency	(Integral QE for distribution centered ~520 nm) 12% for EM section and 7.5% for the Hadronic section
Pulsed photocurrent	Maintain ~5% linearity for 10 ⁶ p.e. in a 10 ns pulse
Average photocurrent	0.5 nA (uniform illumination 13 mm circle)
Lifetime photocurrent	1 mC for 50% gain loss
Uniformity	±15% over 13 mm diameter circle, silvered mesh on photocathode OK
Multiplier	
Gain	6×10^3 nominal
Linear Pulsed Current	25 mA for 10 ns with <2% nonlinearity, 100 mA for 10 ns with <5% nonlinearity
Rise time (20%-80%)	<4 ns (10 mA into 50-ohm load)
Fall time (80%-20%)	<6 ns
Pulse width (FWHM)	<8 ns
Width at baseline (5%)	<12 ns
Pulse shape at high currents	Pulse width (FWHM) changes less than 20% for output pulses between 1 mA and 50 mA
Average anode current	1 μ A
Lifetime anode current	100 C for 50% gain loss
Anode capacitance	<10 pF to all other electrodes combined
Inductance of anode and last dynode	<4 nH
Voltage divider	Simple integer ratios of voltages allow Cockroft-Walton "charge-pump" type base. Typical example: 0 200 400 500 600 700 1000 V for Anode, D6, D5, D4, D2, D1, K.
Noise	
"Single-photoelectron"	<100 KHz at 0.5 of 1 p.e. peak
Afterpulsing	< 0.1% / p.e. for afterpulses >10 p.e.
Pre-pulsing	Arrival < 5 ns early amplitude < 0.5% of signal from photocathode
Stability	
Long term	< 1%/month
Pulse recovery (target)	< 1% gain loss 100 ns after pulse of 10 mA \times 10 ns

The maximum ionizing dose at the design luminosity at EM shower maximum varies from 2.7 krad/year to 6 krad/year for the barrel region and increases up to 570 krad/year at the inner edge of the endcap region ($|\eta| = 3$) (Fig. 6-7). An extensive program studying the damage to scintillators and fibers subjected to radiation has shown that the barrel calorimeter can maintain good performance over the duration of the SDC experiment with the use of commercially available plastics. This R&D program is being continued to find plastics which are even more radiation resistant and which are needed to equip the inner endcap regions where the radiation level is extreme. One difficulty of the radiation damage studies lies in the fact that the radiation damage tests must be accelerated to give useful information in a reasonable time; the dose rates are necessarily much higher than the rates that will be encountered at the SSC.

The results of the various radiation damage studies based on source exposures of individual calorimeter components and electron beam exposures of complete instrumented calorimeter modules are given below. The comparison between results from these two types of tests provides the ingredients for building a model for the interpretation of future source tests of individual tile-fiber assemblies. These are simpler and quicker to perform than beam tests of calorimeters. Future plans to improve the radiation resistance of scintillator and fibers are outlined below as well as the steps necessary to ensure that the assembled calorimeter modules will perform as expected under irradiation. Investigations of the influence of dose rates and of the gas surrounding the scintillator will continue.

6.5.1. Tests on components

A number of radiation damage tests, including bench tests with ^{60}Co γ -sources and low-energy (3 MeV) electron-beam tests, were performed on scintillating tiles and wavelength-shifting (WLS) fibers.

Identical tile/fiber combinations were exposed to ^{60}Co γ -sources at the Phoenix Laboratory in the University of Michigan, the High Energy Physics Laboratory in the University of Tsukuba [25] and the Nuclear Science Center in Louisiana State University. The ratio of light yield before and after irradiation is plotted as a function of total dose in Fig. 6-29a and 6-29b. These results show good agreement with the test beam results. There are no differences in uniformity across the tile surface before and after irradiation up to 5 Mrad. No dose rate dependence of the radiation damage was observed in the range 6 krad/hr to 1 Mrad/hr.




FIG. 6-29. The ratio of light yield after irradiation to before irradiation as a function of total dose using a ^{60}Co source. Data shown are for (a) SCSN81 scintillator tiles read out by Y7 and (b) BCF91 waveshifter fiber.

Tests of the different fibers, tiles, splice joints and machine speeds for groove cutting were carried out at Florida State University using a 3 MeV electron beam as a radiation source. The observed loss of light

agrees with that measured in the ^{60}Co exposures. There are no differences in the uniformity across the tile before and after irradiation. The splice joint between waveshifter and clear fiber did not show any optical or mechanical damage with 1 Mrad of irradiation. Various tests on machine speeds were also performed to obtain a higher quality internal groove surface. The preliminary results indicate that the light loss due to radiation damage can be reduced with higher quality grooves.

6.5.2. Radiation Hardness Tests with assembled EMC modules

A total of 16 EMC modules were built at Fermilab and exposed to electron beams at Beijing, KEK, and Orsay, with energies ranging from 1 to 2.5 GeV. Damage profiles were obtained by scanning the tiles with radioactive sources or with electron beams, and by comparing the light yields before and after irradiation. Total doses, measured at shower maximum with radiochromic films and alanine plastics, reached 4 Mrad. The average dose rates at shower maximum ranged from 0.7 to 20 krad/hr.

Module assembly

To aid in choosing the best available materials, samples of scintillator plates, wavelength shifting fibers, and clear fibers were exposed to a ^{60}Co source at Saclay [28]. Spectral properties, transmission, and light yield of tiles and fibers were studied as a function of radiation dose. Based on the results of these tests, SCSN81 was chosen for the scintillating tiles and BCF91 and Y7 were chosen for the WLS fibers.

Tiles used in the modules were 2.5 mm thick SCSN81 polystyrene-based scintillator manufactured by Kuraray. A 1 mm diameter fiber was embedded in a box-shaped path (Fig. 6-30). For the modules built in the U.S., BCF91 from Bicron was used as the WLS fiber, and the fibers were connected directly to the phototubes. For the modules built at the University of Tsukuba, Y7 from Kuraray was used as the WLS fiber, and the fibers were spliced to clear fibers after exiting from the tile. The edges of the tiles were painted with BC620 reflective coating or polished. The tiles were wrapped with aluminum foil or white paper to optimize the reflectivity of the tile surface.

The light yield from different tile/fiber combinations was measured using UV light sources and radioactive sources (^{137}Cs or ^{106}Ru). Tiles were sorted by light yield, so that tiles with higher yield could be placed near the front of the stack where the higher yield would compensate for the longer fiber length from tile to PMT. The tile-to-tile uniformity was 5 to 10%. After testing for light yield and sorting, the fibers were glued to a cookie. The cookie and fibers were faced off using a diamond cutter. In each module, a set of radiochromic films was imbedded in each layer and alanine dosimeters were placed at EM shower maximum.

The absolute light yield per 2.5 mm thick tile was measured using cosmic ray muons. With a Hamamatsu R329 phototube, a typical tile had about 1.5 p.e./mip. With a Phillips XP2081B phototube, which has higher quantum efficiency in the green wavelength region, the absolute light yield increased to about 3 p.e./mip.

After assembly, all modules were calibrated using a ^{137}Cs source (see below) and sent to three different sites where systematic studies were made: at KEK (lead by the Tsukuba group [29]), Orsay (by the Saclay group [26]) and Beijing (a joint effort of IHEP, FNAL, Purdue, and FSU [27]).

Results of the radiation damage beam tests

Relative calibration was provided at Orsay and Beijing by source tubes. A longitudinal response profile is shown in Fig. 6-31. Transverse scans were also made in order to check that radiation damage does not introduce unacceptable transverse nonuniformities. These are shown in Fig. 6-32 for KEK data. In this case, an absolute calibration was performed using a ^{106}Ru source and data were obtained by scanning an electron beam across the module before and after a 1 Mrad dose. The induced nonuniformity is less than 2% (which is our specification.) This result was confirmed in the Beijing tests for exposures up to 6 Mrad (Fig. 6-33).




FIG. 6-30. Radiation hardness test beam module. The fiber groove path and profile are also shown.

The data from Orsay and Beijing showed that the damage of a tile deep in the stack was small; the relative source calibration was used to determine the damage profile. Different irradiation atmospheres (N_2 and air) were tried, as were different annealing atmospheres (N_2 and air). The best results were obtained in air, as shown by data from Orsay [26].

The longitudinal damage profile is shown in Fig. 6-34 for the Beijing data. Note that the damage profile is slightly broader than the energy deposition profile. The additional damage comes from the WLS fibers being directly irradiated by the electron beam at the edge of the module. It is thought that the BCF91 batch was poor. Furthermore, the WLS fibers were not spliced to clear fibers, so that the radiation damage was increased.

Data from the KEK tests are shown in Fig. 6-35. The results from all three sets of tests are consistent within about $\pm 10\%$.

The beam test results can be summarized as follows:

- The goal of the first round of tests was to prove survivability in the barrel region for 10 years at luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ using commercially available plastics. The maximum dose is 0.6 Mrad for that integrated luminosity. For this absorbed dose, the light loss at shower maximum is about




FIG. 6-31. Longitudinal source response profile of a Beijing module.




FIG. 6-32. Transverse scan using an electron beam for the KEK tests. Data are shown for two strips traversing the tile both before (diamonds) and after (crosses) 0.6 Mrad applied dose.

20–30% and there is good agreement for modules tested at all three sites. As seen in Fig. 6-8a, this loss induces a degradation in energy resolution which is less than the EM resolution of all energies [29]. One can go further and make the degradation negligible, by upgrading the barrel to two EM depths segments (Fig. 6-8b).

- The uniformity over the tile does not change after irradiation up to 6 Mrad, even if the irradiation was not uniform over the tile.
- Modules irradiated in N_2 gas do not show any higher resistance to degradation than the modules exposed in air.
- There is no clear evidence for a dose rate dependence of the radiation damage for rates greater than 10 krad/hr.




FIG. 6-33. Transverse scan using an electron beam for the Beijing data. The relative output (normalized to calorimeter response before irradiation) is shown for cumulative doses (at EM shower maximum) up to 6 Mrad.




FIG. 6-34. Damage profile from source tube scans immediately after irradiation. The data are normalized to no damage at depth $t = 19$.




FIG. 6-35. Light yield ratio after irradiation at shower maximum for KEK modules as a function of total dose at shower maximum.

Future plans

Even though much more work remains to be done to investigate long term, low rate exposures, we can be confident that the scintillating tile/fiber calorimeter will have little degradation in performance in the barrel region from radiation damage. For the endcap region, we plan to allow for replacement of the scintillator every few years in the regions of highest dose (forward-most section of the EM calorimeter). R&D is continuing to produce tiles and fibers that are more resistant to radiation as is discussed below. A technique is also being developed to correct for radiation damage by mapping and masking at the end of the fiber bundle to achieve tile-to-tile uniformity. Higher light yields would help this correction method. The most likely tile/fiber candidates will be exposed in a long term low dose rate irradiation at the University of Michigan (~ 3 Mrad at a rate of < 100 rad/hr). A list of future areas of study includes:

- Wavelength shifting fibers: Higher light yields can be obtained by increasing the fiber density in the tiles.

- **Light guide fibers:** Silica fibers, which have only a few % decrease in attenuation length at 10 Mrad, can be used as light guides in the region of highest dose.
- **Optimizing emission spectrum of waveshifter** to high quantum efficiency regions of photocathode sensitivity.
- **Longer wavelength tile/fiber combinations:** Initial tests show that green tiles (SCSN81 with 60 PPM of Y7) with red WLS fibers (SCSF+3HF+R3) have much improved radiation resistance.

A program of irradiation at Michigan followed by electron beam tests at Beijing will continue in 1993. Initially the options discussed above will be tested. As discussed below, new scintillator, as it becomes available, will be added to the test program.

6.5.3. Radiation-hard scintillator development

Several scintillation-calorimeter based experiments have evaluated the effects of radiation on commercial scintillators[30,31]. The issue of radiation damage to scintillator base polymers and to scintillation fluor systems has been a subject of considerable study in the last few years. An excellent review of this subject is contained in Ref. 32.

The principal light loss mechanism in the base polymer is absorption by color centers formed by oxidation of radiolysis products. Some of these color centers anneal. The annealing rate can depend on the rate of oxygen diffusion into the plastic. It also depends on the temperature and pressure. The color centers initially absorb light at short wavelengths, with the absorption edge moving towards longer wavelengths as the total absorbed dose increases. The obvious solution is to work with scintillation and waveshifted light at longer wavelengths. However, it may be possible to develop plastics which do not form color centers at all, and thus produce plastic bases which are intrinsically radiation-hard.

The scintillation fluors themselves can also be destroyed by radiation. However, many dyes show great resistance to such damage. Therefore, although some care must be taken to avoid colored radiolysis products of the dyes themselves, the principal requirement on the dye system is that the energy be efficiently coupled to the photon detector at as long a wavelength as is possible.

Development Program

The objective of the program is to obtain a plastic scintillator with high resistance to radiation doses in excess of 10 Mrad, light yield comparable to the highest achievable today (in PVT), mechanical properties appropriate for high-volume/low-cost fabrication of tile/fiber assemblies, and long-term resistance to aging.

The specific areas under investigation comprise :

- The development of new plastic bases (Bicron "RHxx") which are intrinsically more radiation hard than present day styrene or vinyltoluene based systems (Fig. 6-36). Scintillating fibers with these new bases exhibit significant resistance to radiation damage [33] (Fig. 6-37). One of our major goals is to improve this technology and to use it for the fabrication of scintillator plates. At present, plates using this material show reduced radiation damage, but they exhibit crazing on very short time scales. Work is currently in progress to improve the radiation hardness and the mechanical resistance of this system.
- An evaluation of the effects on radiation resistance of the polymerization cycle and subsequent processing of the sheet plastic.
- An evaluation of candidate additives and stabilizers (both from commercially available and proprietary sources).
- An evaluation of effects of impurities introduced into the plastic base during the fabrication process (such as from mold release agents).
- Optimization of the full fluor system (primary dye, shifter and wavelength shifter), including optimization for speed, emission/absorption overlap, and the final match to commercially available photocathodes.




FIG. 6-36. Light yield of RH-1 scintillator plate as a function of path length and time following irradiation to an absorbed dose of 1 Mrad




FIG. 6-37. Light yield of RH-I based fibers irradiated to 3 Mrad.

Evaluation of the plastics produced by Bicorn Corporation under this R&D program has been carried out at Florida State University, Louisiana State University and Argonne National Laboratory. As an indication of the success in this effort, Fig. 6-38 shows results from two recently produced 2-inch diameter test disks which were irradiated at 600 krad/hr for 4 hours using a ⁶⁰Co source [34]. The signal after 2.4 Mrad of radiation recovers to essentially 100% of the initial light output. Although this is one of the best results obtained to date, subsequent tests of similar materials have confirmed recovery to between 90 and 95%.




FIG. 6-38. Yield ratio from 3 test disks of material described in the text, read out using 2 wraps of 1 mm BCF91A fiber, after 2.4 Mrad irradiation.

6.6. Mechanical and integration issues

6.6.1. Introduction

The barrel and endcap components of the central calorimeter assembly each consist of electromagnetic and hadronic subcomponents. The steel and lead assemblies of the barrel calorimeter are part of a single wedge unit. The endcap has a removable EM section and at small angles a removable hadron absorber plug. Although the barrel wedge and endcap components are straightforward in design, their assembly and support present formidable problems. The support system for the calorimeter and other parts of the detector system attached to the calorimeter must be minimal to allow space for readout and access.

Mechanical designs for the forward calorimeter are being developed for each option as will be discussed in Section 6.10. Option-independent issues include integration with the forward muon system and accelerator, as well as radioactivation.

6.6.2. Barrel calorimeter

The barrel calorimeter (Fig. 6-39) is divided into two rings, each with 32 keystone-shaped wedges. The calorimeter is constructed with 0.05×0.05 segmentation in EM, and 0.1×0.1 segmentation in hadron sections. Each wedge covers from $\eta = 0$ to $|\eta| = 1.4$, spanning 4 EM towers in ϕ , and 28 EM towers in η . Each hadron calorimeter tower contains iron/scintillator HAC1 and HAC2 sections. The calorimeter is assembled out of pairs of half-wedges. This allows the readout fibers to be routed up the sides of each tower, and simplifies the routing of the fibers at each layer of the calorimeter. The two half-wedges are then bolted together into a full wedge which weighs about 37 tons, and can be transported, manipulated by cranes of reasonable capacity, and individually exposed to a test beam. Each wedge carries its own phototubes and electronics, and arrives at the SSCL as a tested and calibrated module.

At the SSCL the wedges are stacked into rings on a support cradle. During and after stacking, the mechanical strength of the ring is provided by the iron structure of the hadron calorimeter. The load is carried by the support cradle and attachment plates bolted onto the ends of the wedges at their inner and outer radii. An analysis of the forces encountered during stacking, and the sizing of the support arms and attachment plates has been made and is currently being refined.

Barrel hadron calorimeter

The barrel hadron calorimeter modules are keystone-shaped blocks of iron containing a pattern of "staggered slots" to contain the scintillator tiles. The staggering of the slot pattern allows mechanical loads and magnetic fields to bypass the scintillator tiles, while still maintaining appropriate sampling in projective towers.

The absorber assembly method uses the conventional technique of stacking iron plates in a plane parallel to the scintillator plates, to build up the wedge structure. In this technique the scintillator slots are produced by milling channels (6.35 mm deep by 15 cm typical width) into each iron plate. Grooves are milled the length of each plate to accommodate the source tubes. In our present specification, each plate is processed to a flatness tolerance of ± 0.33 mm to ensure an appropriate stacking tolerance (3.1 mm statistical, 28.1 mm worst case) on the completed stack. The stack is assembled by plug welding each layer together. A stacking fixture guarantees that the stack stays aligned and that the interior edge of the half-wedge stays perpendicular to the stack to an accuracy of ~ 1.3 mm. After the stacking is completed, a final machining pass is made on the outer surfaces of the block. This machining generates the readout channels for the fiber bundles and determines the keystone shape of the finished block to an overall accuracy of 0.8 mm.

The block is then cleaned and deburred. Source tubes are inserted into the grooves. Prior to final assembly of the calorimeter, every tile is accessible to sources. To allow for rapid periodic monitoring, source tubes sampling two depth layers in each η/ϕ segment for each longitudinal section (EM1, EM2, HAC1, HAC2) are routed to the outside of the barrel. These special tubes in the steel require a curved

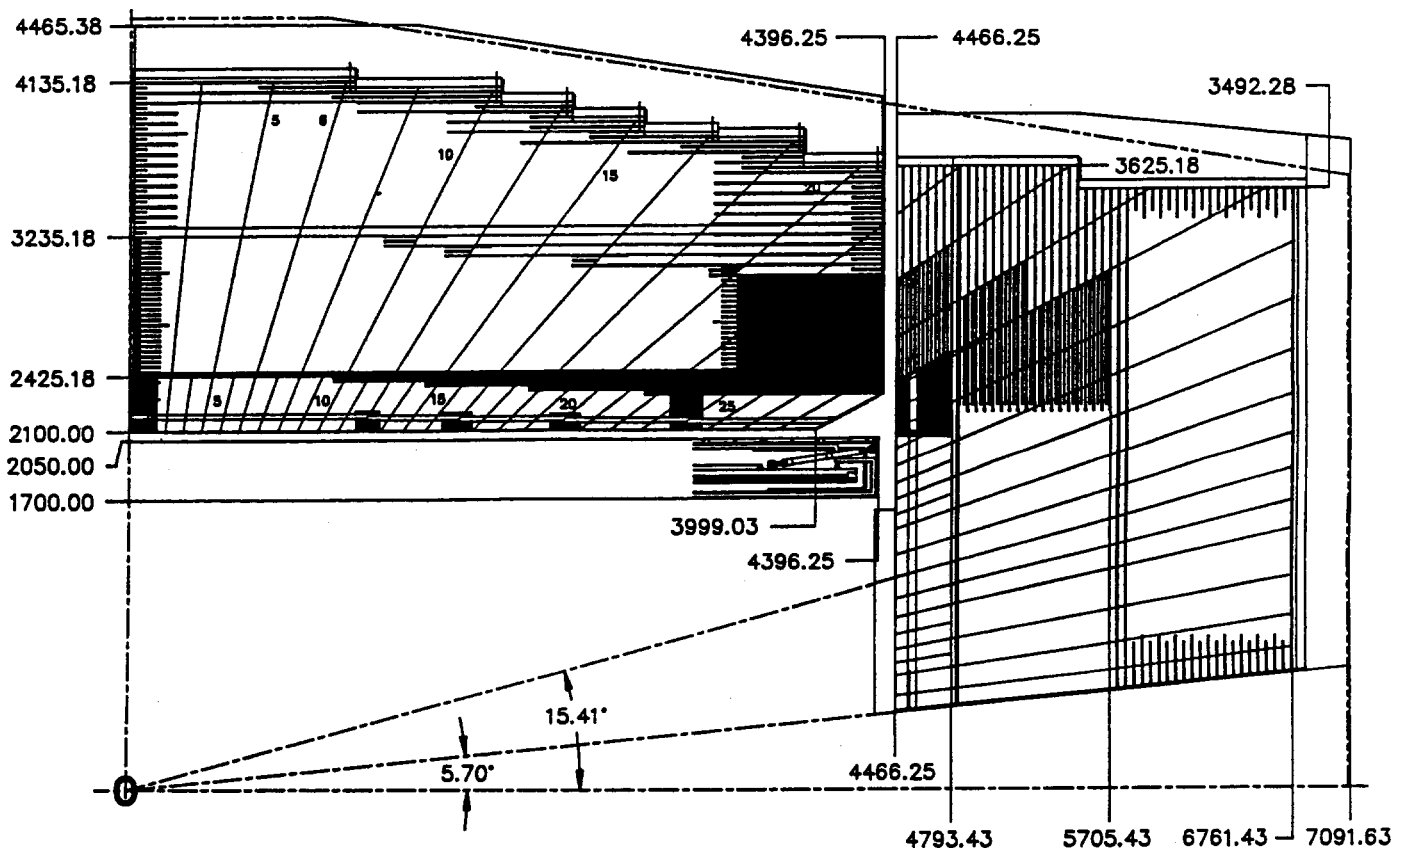


FIG. 6-39. Barrel Calorimeter—a quarter cross section of the tile/fiber central calorimeter. The EM sections are fabricated with cast lead alloy; the hadron sections are formed by stacked steel plates welded together. A shower maximum detector is included in the EM section.

groove to be milled for the tube to come out parallel to the wedge face. Occasional access to all other source tubes to allow monitoring of every tile can be achieved by pulling back the endcap.

The scintillating tiles are inserted into the absorber slots, with each tile being held firmly against the slot surface and source tube by simple spring-metal clips. The current model calls for installing complete tile/fiber assemblies. A possible option being investigated is to insert the fibers after tile insertion, to reduce the difficulty associated with handling tiles with long fibers attached.

The readout fibers are carefully routed in a single layer along staggered channels on the surface of the wedge. The minimum depth of the groove is determined by the outside diameter of the splice ferrule on the fiber. This procedure has been tested in the prototype calorimeters for the test beam.

Barrel electromagnetic calorimeter

The barrel EM section is made of cast lead. The lead is held by a stainless steel frame with a flat plate at the outer radius and a plate at the inner radius that is formed to provide end supports as well. The remainder of the frame consists of thin 0.5 mm welded stainless steel bulkheads that constitute the scintillator tower boundaries in η . Into this frame are cast a series of parallel lead plates using the frame as the supporting structure. The result is a half wedge with slots for containing the individual scintillator tiles. There are 64 of these units in each half barrel of the central calorimeter. Each unit contains 2×28 projective towers of scintillator tiles. These towers are matched to the hadronic section of the barrel when the hadronic and electromagnetic units are mated. In order to compensate for the increased depth of the absorber plates as $|\eta|$ increases, the hadronic section of the calorimeter is allowed to intrude into the EM




FIG. 6-40. Barrel calorimeter EM section. Tower boundaries are structural steel bulkheads.

section in progressively deeper steps as $|\eta|$ increases. A cross section of the barrel EM calorimeter is shown in Fig. 6-40.

Finite element analysis done on the barrel EM section shows that the design features are well within the limits of feasibility [35]. Further refinement will be pursued to optimize the structural integrity of these units. We expect to use a lead alloy (commonly used for battery plates) to provide mechanical strength and resistance to creep not available from pure lead.

6.6.3. Endcap calorimeter

The endcap calorimeter covers the range from $\eta = 1.4$ to $\eta = 3$. Each endcap is made up of 32 wedges that cover and close one end of the barrel calorimeter (Fig. 6-41). The endcap has four depth segments; two EM sections, followed by a first and second hadronic sections. The electromagnetic sections are cast lead with the plates oriented perpendicular to the beam axis and are designed as a separate detachable unit to provide the possibility of removing and replenishing the scintillator tiles if they are damaged by radiation. The hadronic sections of the endcap calorimeter are constructed of steel plates, separated by staggered spacers, that are also perpendicular to the beam axis. A cross section of this construction is represented in Fig. 6-42.

Finite element analysis has been done on the individual sections of the endcap and is currently being done on the entire endcap assembly. The analysis of the endcap EM section has determined the stresses in the structural elements and in the lead absorber plates. The combined stresses range from 300 psi to 4000 psi in the supporting structural plates, and from 25 psi to 75 psi in the lead absorber plates. These are acceptable numbers for the proposed lead alloy.

The removable EM section has 32 wedges per endcap which contain all the EM towers plus other associated elements of hadron towers continued from the barrel (see Fig. 6-42). Thus, when tiles in the endcap need to be replaced due to radiation damage, the endcaps will be rolled back one meter so the wedges can be removed one by one. This is a major operation which will be done only after several years




FIG. 6-41. Endcap Calorimeter wedge. The azimuthal segments of the endcap calorimeter are formed as separately mounted wedges to facilitate removal if refurbishment is needed as a result of radiation damage.

of operation.

6.6.4. Removable hadronic plug

The hadronic section scintillating tiles for $|\eta| > 2$ will also require periodic replacement. Figure 6-43 shows our baseline design of a removable hadronic plug, which is supported from the back of the endcap. The absorber structure of the monolithic plug consists of vertical layers of annular steel plates, with 50% of the area in ϕ covered by trapezoidal shaped steel wedge plates. Between the wedge plates of each two layer cell are "pizza pans" containing tiles which are installed from the outer surface of the plug. Each of the 32 ϕ sectors in the hadronic plug contains 8 towers with a cell structure matching that in the surrounding endcap. Efforts have been made to keep the projective endcap/plug η -gap thin. The annular absorber plates and wedge plates of each cell are fastened together by arrays of short, welded steel tie-rods. This simple construction allows easy access to the pizza pans and provides a rigid structure to resist gravitational and magnetic loads (about 300 tons).

Readout fibers are routed radially outward from the tiles, and follow the plug's outer η boundary to the rear of the plug where the PMT's are mounted. Source tubes, on the other hand, exit the plug's inner $\eta = 3$ boundary, and are routed to the back plate where source drivers are located.

The size of the removable plug is determined by the expected radiation damage to tiles and by the ease




FIG. 6-42. Endcap calorimeter transverse section. The components are as in the barrel, Fig. 6-39. The EM and HAC transverse tower structures are shown.

of plug replacement. The damage and radioactivation in the absorber will determine how frequently plug components need to be replaced. The baseline plug, Fig. 6-43, weighs about 100 tons and contains 7680 tiles. The plug's high radiation damage region is limited to the front, HAC1 section, beyond $\eta = 2.4$. We are currently investigating smaller plug configurations with fewer towers.

6.6.5. Barrel support system

The basic design criteria for the SDC calorimeter support structure is that it support the calorimeter components safely under several different loading conditions, producing only minimal impact on the muon system. The calorimeter support system must transfer the loads from the calorimeter to the muon steel.

Assembly of the muon chambers must therefore be considered in the design of the calorimeter support. Access under the calorimeter is a difficult issue at best and much thought has gone into the design of this support to create sufficient access to both muon and calorimeter electronics. Starting from the muon steel and progressing up to the calorimeter wedges the first section of the support is composed of massive steel "picnic tables" which are bolted to the muon steel. These are the load transfer parts which hold the calorimeter above the lowest muon chamber. They weigh about 20 tons each and there are about 25 of them placed end-to-end to create the continuous rail that the calorimeter rides on. During installation assembly, these tables are roughly positioned and bolted down, then adjustable plates are levelled and grouted in place. These adjustable plates form the rolling surface for the next layer, the Hillman rollers. The calorimeter uses 500-ton rollers. They are equipped with side mounted rollers to prevent the load-carrying rollers from falling off the rails.




FIG. 6-43. Two views of the removable hadronic plug. The structure consists of annular absorber plates and trapezoidal wedge plates connected by arrays of short tie-rods, welded to every other plate. Pizza pans containing tiles for eight towers are removed radially from absorber slots when the cantilever supported plug has been withdrawn from the endcap. Readout fibers are routed under the thin, sheet metal panels to PMTs at the back plate. Source tubes are under the bore skin. Pizza pans and tile/fiber units are installed using a "last in = first out" assembly sequence which simplifies replacement of those most subject to radiation damage.

Above the rollers are elastic plates or pads made of a neoprene composite. Above the elastomer pads are hydraulic cylinders capable of adjusting the vertical height of the calorimeter and, used in pairs, adjusting roll and pitch. Another technique borrowed from CDF is to connect the different pairs of cylinders to change the four point support to a three point support. By joining two cylinders and keeping the other two independent, the loads can be evenly distributed in an analyzable way. The vertical hydraulic cylinders are designed with a locking nut, so they can be fixed after adjustment and the pressure bled off.

The three-point support works fine while the hydraulic pressure is supplied; even if the floor were to shift and cause deflections of the muon steel, it could adjust to the new shape. However, during most of the calorimeter's life it will be standing on the fixed locknuts. This means that a careful analysis of the torsional rigidity of the calorimeter must be performed to determine what happens to individual cylinders and rollers when the rails deflect and change the load distribution on the the locked cylinders. If the calorimeter is determined to be torsionally rigid, the elastomers will have to be designed to take large displacements at a relatively constant load. If the calorimeter has minimal torsional rigidity, it will follow the deflections of the support rails and distribute its load to the rollers and hydraulics more evenly. This analysis is in progress now.

Horizontal adjustment capability is built into the design. Horizontally-acting hydraulic cylinders will be attached to the Hillman rollers and push against a sliding plate holding the vertical cylinders. The controller to make sure all the cylinders are behaving properly is being defined now.

Above the vertical hydraulic cylinders is the cradle structure; two plates support each half-barrel and end cap. These plates are connected to each other by a longitudinal beam to absorb forces parallel to the center-line of the calorimeter. These plates weigh approximately 12.5 tons and are made of A36 structural steel. They and the longitudinal beams which connect them must be designed to allow access to the electronics crates below the calorimeter and the removal of cards from the crates.

The wedges are attached to the cradles by an interface piece which allows access for bolts going into both pieces. The interface pieces also create spaces where cables and utilities can cross through the cradles. Finally, the entire calorimeter and support cradle is moved on the rail platforms by horizontally acting hydraulic cylinders. These hydraulic cylinders will probably be connected to the Hillman rollers on one end and the top rolling rail at the other, but these connections have yet to be designed.

6.6.6. Endcap support system

The end cap support system is patterned after the barrel support system. Similar components were chosen for it, and the operating principle is the same.

6.6.7. Forward calorimeter integration and SSC interface aspects

The successful realization of the seamless transition between the endcap and forward calorimeter will require attention to some details. The skin of the endcap and removable plug calorimeters on the inner conical surface at $\eta = 3$ should be as thin as practical; there should be no major structural elements there.

The forward calorimeter will require careful integration with the forward muon system. It will be mounted on the inside of the forward muon toroid steel. Moreover, the design of the iron shield required on the inner bore of the forward muon system is of mutual concern. Its function is to protect forward muon wire chambers from the charged particle and neutron flux from the front face of the forward calorimeter, and from the beam pipe. This shield must not interfere with the transition from the central calorimeters to forward calorimeter at $\eta = 3$. The front face of the forward calorimeter (including the extra 1λ radial allowance) must be unobstructed.

The principal SSC machine elements affected by the location of the forward calorimeter are the first quadrupole and the shield needed to protect it from the energy of those interaction products not intercepted by the forward calorimeter. Since the best performance is obtained with the forward calorimeter as far from the interaction point as possible, the shield will become an integral part of the forward calorimeter design and construction.

Radioactivation and access.

Radioactivation of forward calorimeter elements in the region nearest the beam pipe will be similar to that associated with near-target environments in fixed target machines, and care must be taken in planning forward calorimeter service activities.

6.7. Shower-maximum detector and electron identification

6.7.1. Electron and photon identification

The SDC detector design is strongly influenced by the major goal of identifying electrons and photons.

Electron identification

Electrons can be “isolated,” “partially-isolated,” or “nonisolated.” The benchmarks for isolated electrons are W and Z electronic decays. New physics, such as $Z' \rightarrow e^+e^-$, also produces isolated electrons. The benchmark for partially-isolated electrons is b decay, $b \rightarrow ce\nu$, and especially the top decay sequence $t \rightarrow b \rightarrow e$. Non-isolated electrons, such as electrons from c decay, are backgrounds to physics of interest. Another important source of background electrons is the conversion of photons into electron pairs, especially in the beam pipe and in the first layers of the inner silicon tracker.

Some selection criteria are common to all sources of electrons. Electrons form characteristically narrow showers in the EM section, with little leakage into HAC1 section. The momentum measured in the tracking system should be equal to the energy measured in the EM. The shower position can be precisely measured in the Shower Maximum Detector (SMD), which consists of crossed scintillator strips incorporated into one gap of the EM calorimeter [36]. The shower position measured in the SMD should agree with the extrapolated position of the electron track. The transverse profile of the EM shower can also be studied with the SMD. Further information can be derived from the massless gap, which consists of the first scintillator tile of the EM calorimeter read out as a separate electronic channel. The shower development can be followed from the massless gap, through the EM and the SMD to the HAC. All these elements of information will be used to select electron candidates.

Isolated electrons are more likely to signal new physics than nonisolated electrons. Isolation criteria involve rejecting events if the EM or HAC towers near the electron candidate have energies comparable to the energy deposited by the electron candidate or at least greater than some minimum value. Mini-isolation cuts based on the SMD can also be used to enhance the signal for partially-isolated electrons.

The data necessary to implement these selection criteria are not all available in time for the Level 1 trigger. EM and HAC calorimeter data are available, but only limited information is available from the outer tracker and from the SMD. Track segments from the outer tracker provide a loose momentum cut at Level 1 ($p_t > 10$ GeV/c). The SMD provides matching information in ϕ , with a very loose energy threshold.

The selection criteria for single electrons at Level 1 require an EM signal with, for example, $E_t > 20$ GeV with a corresponding stiff track and a loose track-shower match. The HAC/EM ratio can be required to be small, but isolation cuts are not imposed at Level 1. Multiple leptons can be selected at Level 1 with looser cuts than those used for single electrons or muons.

At Level 2, more information is available from the tracker (including the inner silicon tracker) and from the SMD. Photon conversions can be found and eliminated at Level 2 using the silicon tracker. Once the silicon tracker has eliminated the photon conversions, fake electron candidates arising from $\pi^\pm\pi^0$ overlaps can be eliminated by using the momentum-energy match (tracker-EM), and the improved track-shower match with the full SMD granularity.

At Level 3, the SMD can be used to identify electrons that have lost significant energy through bremsstrahlung. Detailed correlations of the data from the trackers, the massless gap, the calorimeters, and the SMD can also be performed at Level 3.

Photon identification

Photons of interest for new physics searches will usually be isolated. The benchmark is Higgs decay into $\gamma\gamma$ for neutral Higgs with mass below 130 GeV, with isolated leptons from associated W decays to enhance the signal.

Background to isolated γ 's comes from π^0 decay. Single γ 's can easily be distinguished from γ pairs from π^0 decay in the SMD if the shower peaks from the two γ 's are separated by at least two low bins (γ separation greater than twice the rms width of the showers). Monte Carlo studies have shown that we can reject 80% of the π^0 's at 25 GeV with a γ efficiency of 80%, based on the shape of the energy distribution in the SMD. These studies have also shown that separation could be achieved at higher energies, by using narrower SMD strips, because there is a narrow core in the energy distribution of EM showers. We are

studying the need for better π^0 identification. The separation of γ 's and π^0 's can be studied at even higher γ energies on a statistical basis using data from the massless gap.

6.7.2. Requirements for the SMD

The physics requirements for the SMD were summarized in Section 6.2. The SMD should measure the position and the profile of showers at the maximum of the EM shower. Requirements on the SMD spatial resolution and granularity, and on the SMD energy range and energy resolution are given below. The contribution of the SMD to the different levels of the SDC trigger is discussed above and in Section 8.2. Finally, the effects of radiation damage and the consequences of our choice for the depth of the SMD are presented below.

Spatial resolution and granularity

The SMD is used to match EM showers in the calorimeter to high- p_t tracks in the tracking system, and to separate single EM showers from multiple EM showers and from pions interacting early in the EM calorimeter. Fine granularity is required to measure the center-of-gravity and the profile of the EM showers.

The RMS width of the energy deposition at $6 X_0$ is 1.7 cm for electrons and 3.5 cm for early-interacting pions. The SMD strip width should be smaller than the size of the energy cluster, in order to measure the position and the profile of the cluster. We have chosen 8 strips per EM tower, giving strips which are 1.2 cm wide at $\eta = 0$, or 40% of the 3-cm Moliere radius in the SDC lead/scintillator EM. This is the same ratio as that of the CDF SMD (1.45-cm strips, with a Moliere radius of 3.53 cm). We have two layers of SMD strips to measure the η and ϕ coordinates of the EM showers, with the same segmentation. This choice of 8 strips per tower ensures convenient and economic electronics packaging.

The two layers of SMD strips are placed in the same slot in the EM calorimeter to ensure a close correlation between pulse-heights for matching in the case of partially-isolated electrons. The pulse heights from the two layers are equal within 20% at $E_t = 10$ GeV and within 5% above $E_t = 100$ GeV (at $\eta = 0$). This correlation is especially useful for tagging conversion electrons, where, near 90° , we expect to see a single shower in η , but two showers in ϕ .

The SMD will cover the full η range of the barrel and endcap calorimetry. The fiducial volume for electron acceptance will be determined in part by the SMD coverage. Losses due to cracks or dead spaces should be below 5%.

Length of SMD strips

Ideally, the SMD strips would be the same length as the EM towers. However, given the existing financial constraints, we have decided to gang readout fibers together, making logical strips which are longer than the EM towers, to reduce the number of electronic channels.

The length of the logical strips is limited by the occupation level of the strip under expected experimental conditions. The SMD should be able to identify electrons that are only partially isolated. A typical example would be the electrons from the decay sequence $t \rightarrow b \rightarrow e$. For a top mass of 150 GeV, approximately 36% of the electrons from the above decay sequence have additional particles with energies above 5 GeV in the same trigger tower (0.10×0.10). Only about 3% have 5 GeV additional energy incident on a single SMD strip of the same length ($0.10 \times 0.05/8$). With crossed strips, these electrons can be identified with good efficiency even if the SMD strips are twice as long as the trigger towers ($0.20 \times 0.05/8$).

Other requirements favor the shorter SMD strips. One of these requirements is the occupancy due to minimum-bias background which depends on the luminosity. A good match would be to limit the SMD occupancy to no more than the occupancy of the EM towers themselves. An EM shower leaves most of its energy on three or four of our SMD strips, so this would limit the length of the SMD strips to two EM towers ($0.10 \times 0.05/8$).

The shorter SMD strips ($0.10 \times 0.05/8$) would also greatly simplify the implementation of the Level 2 trigger. In the current design, the trigger towers contain four EM towers (one hadron tower 0.10×0.10), so that SMD strips two EM towers long would match this trigger-tower segmentation.

Due to financial constraints, we do not intend to implement this preferred segmentation in the SMD electronics for the first period of SSC operation, but it will be retained as an upgrade path. In fact, because of mechanical constraints, the physical strips will be only one tower long. Initially, four readout fibers will be brought to each photo-detector channel. Thus the detector will be mechanically divided into strips measuring $(0.05 \times 0.05/8)$, while each electronic channel will measure $(0.20 \times 0.05/8)$. When additional funds are available, the number of electronic channels can be increased.

Energy resolution and dynamic range

The SMD should measure EM showers for transverse energies from 10 GeV to about 1 TeV. The lower limit is set by the momentum spectrum of electrons from the top decays $t \rightarrow b \rightarrow e$. The SMD information is most important at the lower values of E_t where the backgrounds are higher. Only minimum energy resolution is required, because the spread in the energy deposit in the SMD is about 25% due to sampling. The relative size of the energy deposit in neighboring strips must be known to about 10% to avoid degradation of the position measurement.

The depth of the SMD is chosen so as to ensure maximum efficiency for electrons with $E_t = 10$ GeV. Figure 6-44 shows the energy deposit for different incident electron energies as a function of depth in radiation lengths (X_0). Shower maximum is at $6 X_0$ for a 10 GeV electron, and at $11 X_0$ for a 1 TeV electron. At $6 X_0$, the energy deposit from a 1 TeV electron is only 33 times larger than that from a 10 GeV electron. With the SMD at $6 X_0$ at $\eta = 0$ in the barrel calorimeter, we obtain maximum efficiency for the low-energy electrons and a reduction in dynamic range by a factor of three. At the end of the barrel, shower maximum for $E_t = 10$ GeV is at $7 X_0$. We place the SMD in the barrel EM such that the total number of radiation lengths, including those from the magnet coil, varies from six at $\eta = 0$ to seven at $\eta = 1.4$.

For $E_t = 10$ GeV, shower maximum varies from about $7 X_0$ at the outer edge of the endcap ($\eta = 1.4$) to about $9 X_0$ at the inner radius ($\eta = 3.0$). We place the SMD at $7 X_0$ throughout the endcap calorimeter, in order to ensure a smooth transition from the barrel EM.

The location of the SMD makes it vulnerable to radiation damage. Figure 6-45 shows radiation damage as a function of shower depth. The radiation damage is about 3 times worse at $6 X_0$ than at $11 X_0$. The effect of this choice on the survivability of the SMD will be discussed below.

Table 6-10
Dynamic range factors for SMD readout.

Consideration	Factor
Energy range 10 GeV to 1 TeV ($6 X_0$)	33
Downward fluctuation at 10 GeV	2
Upward fluctuation at 1 TeV	1.5
Shaping and center-of-gravity	10
Pedestal separation	4
Total minimum range	4000

The factors determining the required dynamic range of the electronics at $\eta = 0$ are listed in Table 6-10. This range just fits onto a 12-bit ADC, with minimal separation between the pedestal and the useful signal. This choice is justified if there is no noise in the photodetector system; this is indeed the case if the




FIG. 6-44. Energy deposit vs depth for different electron energies. The data were measured to 15 GeV, then extrapolated to higher energies. The curves are taken from Muller [37].




FIG. 6-45. Radiation damage as a function of depth in radiation lengths for different doses (the doses are measured at shower maximum). These data were obtained at Orsay with one of the radiation-damage modules described earlier in this chapter and in Ref. 38.

photodetector is a multi-channel PMT (MCPMT). We have made an effort to fit the SMD signals into a 12-bit dynamic range because this is the accuracy of the switched capacity arrays (SCA) developed as a calorimeter readout option.

A 10 GeV electron at $\eta = 0$ will leave on the average an 80 mip-equivalent signal in each of the SMD layers. The largest signal in a single strip will be about 40 mips. The least count for a MCPMT located at $\eta = 0$ will correspond to half a mip. We hope to obtain two photoelectrons per mip with the 4 mm thick SMD strips, or one photoelectron per count.

Radiation damage

The SMD should be sufficiently robust to survive the radiation doses acceptable to the EM calorimeter and still satisfy the requirements outlined above. The main requirement is to preserve the measurement of the position of the EM showers. A precise energy measurement is unnecessary, since the SMD slot contains a normal tile in addition to the SMD strips. Only relative pulse heights between neighboring strips are required for the essential position measurement; absolute value is used for longitudinal shower development and to reject anomalous EM measurements in the barrel, where there is a single EM depth segment.

The damage in the scintillator is mainly due to a loss of light transmission [39]. Since the path length for the light to reach the WLS fiber is smaller in a SMD strip than in a calorimeter tile, the corresponding signal loss will be smaller. On the other hand, the fiber routing is more complicated for the SMD than for the tiles, so that the loss from radiation damage to the fibers may be greater.

The choice of depth for the SMD represents a compromise between radiation damage on the one hand, and shower profile and dynamic range on the other. Radiation damage is mainly due to EM showers of low energy (~ 1 GeV) and is greatest near $4 X_0$. Placing the SMD deeper inside the calorimeter ($\sim 10 X_0$) would considerably reduce the damage, but the EM showers would be broader, and the two-shower resolution would be degraded. Furthermore, the energy range accessible to a 12-bit readout system is reduced for the deeper SMD, as explained above. We chose $6 X_0$ at $\eta = 0$, varying out to $7 X_0$ at $\eta = 1.4$, to get as much light as possible for electrons with 10 GeV E_t . We chose $7 X_0$ in the endcaps to provide a smooth transition from the barrel region. More light loss can be accepted in the forward regions of the SMD without seriously compromising the position resolution, because a fixed E_t produces up to 5 times as much light as at $\eta = 0$.

We plan to conduct radiation damage tests on SMD strips and assembled detector elements with ^{60}Co sources. These elements will also be incorporated into the new series of calorimeter test modules to be irradiated in electron beams.

6.7.3. SMD beam test results

We have studied the performance of our design with different SMD tile configurations in the MP test beam at FNAL (Section 6.11). The prototype cast lead calorimeter developed at ANL was equipped with different SMD prototypes in separate towers. In the design of the calorimeter lead casting, a space of 9 mm was left for the insertion of an SMD, with or without a standard EM tile.

The following SMD configurations were tested:

- 1) Eight η -strips per tower, 2.5 mm thick.
- 2) Four η - and four ϕ -strips per tower, 2.5 mm thick. The fibers in the ϕ -strips made a 90° turn within the strip to test routing in a very limited space.
- 3) Eight η -strips per tower, 5 mm thick.
- 4) Strips of two widths with highly beveled edges. These are prototypes for the η -strips at the end of the barrel EM.
- 5) Eight η -strips cut from a single piece of scintillator 2.5 mm thick. This prototype tests the cross-talk for strips fabricated using an inexpensive technique.
- 6) Eight η -strips with the WLS fibers along the sides of the strips.

The SMD strips were placed in the calorimeter towers at a depth of $6 X_0$, and were read-out with wavelength-shifting fibers connected with an optical inter-connect to the photodetector. Two types of photodetector were tested: the Philips 64-channel MCPMT and APD's from EG&G.

Data were taken with electrons and pions of energies between 15 GeV and 150 GeV. The electron energy deposits measured in the SMD are shown in Fig. 6-46 for different beam energies. Pulse heights obtained for pions and for electrons at 35 GeV are shown in Fig. 6-47; a pulse-height cut selecting 95% of the electrons retains 5% of the pions.




FIG. 6-46. Energy deposit in the SMD as a function of electron energy. The SMD prototype (8 η -strips 2.5 mm thick per tower) was located at $6 X_0$. It was read out using a Philips 64-channel MCPMT, followed by a preamplifier and a 16-bit ADC. Pulse height distributions (in ADC counts) are given for 15, 20, 27, and 35 GeV electrons.

The rms position resolution was measured to be 2.4 mm (Fig. 6-48). Monte-Carlo simulations show a 2 mm resolution. Systematic effects observed in the beam-test results lead us to believe that we can reach this value.

The profile of the EM showers was measured at different depths using a prototype preshower detector built with 1 mm scintillating fibers. The full width of the showers at a quarter of the maximum was measured to be 1.7 cm for electrons at $6 X_0$. Monte Carlo simulations, as well as the experimental data, show that the showers have a narrow core a few millimeters wide. This narrow core is evident in the test beam data shown in Fig. 6-49.

The same prototype preshower detector was used to estimate the effectiveness of the massless gap to correct for degradation in energy resolution caused by material in front of the electromagnetic calorimeter. With two radiation lengths upstream of the EM calorimeter, the energy resolution at 15 GeV is degraded from 5.0% to 8.0%, and can be restored to 6.3% with the massless gap signal. At higher energies, the degradation is much smaller and the massless gap correction almost completely removes it.

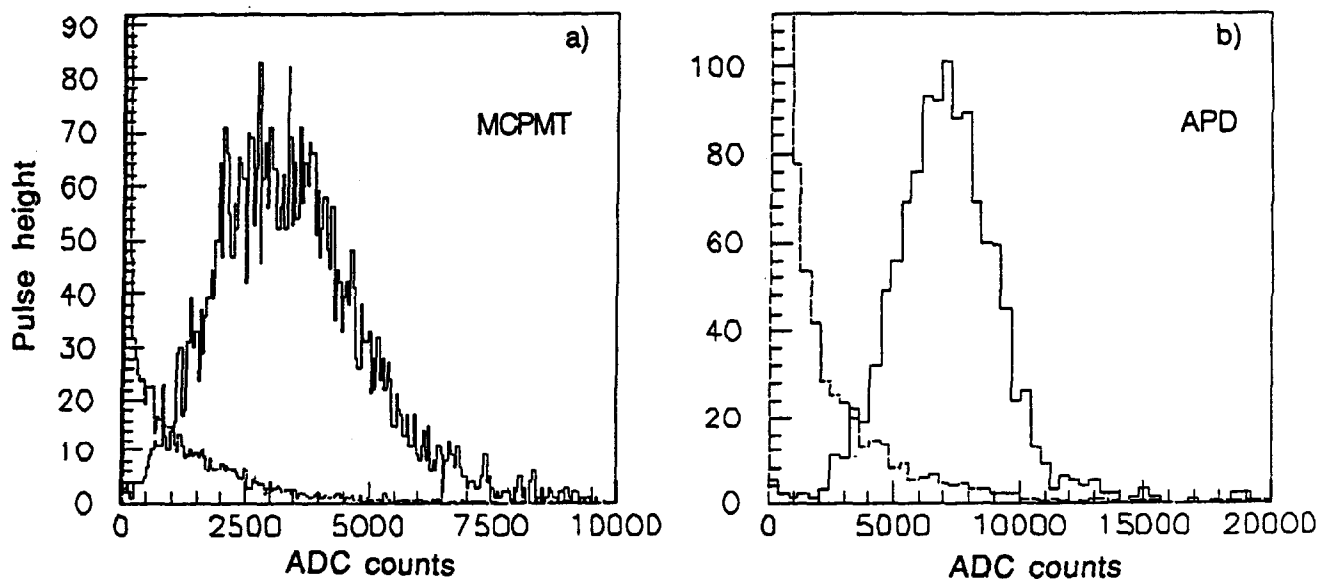


FIG. 6-47. Pulse heights measured with the SMD with 35 GeV pions (dotted) and electrons (solid); a) shows data taken with the MCPMT readout and b) shows data collected with the APD readout.




FIG. 6-48. Position resolution for electrons in the SMD collected with the MCPMT readout and 5 mm thick strips. The difference between the position in the SMD and the measured trajectory of the incident electron is shown. The shower maximum position resolution found is ± 2.4 mm.




FIG. 6-49. Shower profiles at shower maximum for 35 GeV electrons; a) shows data collected with the prototype preshower detector with 1 mm fibers, and the narrow core of the non-gaussian shower profile is evident. The dotted and the solid lines are the measured shower profiles at 5 and 7 X_0 respectively. b) shows the shower profile measured with the SMD (12.75 mm strips, 2.5 mm thick).




FIG. 6-50. SMD segmentation in the barrel calorimeter. The physical strips are shown schematically in the upper figure; the granularity of the electronics is shown in the lower figure. The physical strips are one EM-tower long and 1/8th of a tower wide. The η -strips (measuring η) are shown at the left end of the barrel wedge; the ϕ -strips, at the right end. The electronic channels are shown in the lower figure ganging four physical channels together.

6.7.4. SMD design

Barrel strips

The SMD in the barrel calorimeter will be located in slots in each EM wedge at a total depth (coil included) of about $6 X_0$. The SMD slots contain SMD strips covering four EM towers (one hadronic tower). These slots are $\Delta\eta = 0.1$ long. There are bulkheads in the lead casting at every other η boundary in the EM wedge. The SMD slot contains one layer of normal EM sigma tiles followed by the two crossed layers of SMD strips to measure the η and ϕ coordinates of the EM showers. An extra-thick 26 mm slot is necessary for this collection of scintillator.

The physical SMD strips are shown in Fig. 6-50; they measure $d\eta \times d\phi = 0.05 \times 0.05/8$ and $0.05/8 \times 0.05$ for the two layers of strips. Each strip has a single WLS fiber along the strip length in a slot in one face of the strip. The strips at larger values of η have beveled edges because the bulkheads delimiting the SMD slots are at angles deviating as much as 60° from a right angle. The form of the extreme strips is shown in Fig. 6-51.



FIG. 6-51. Beveled strips for the SMD. An η strip for the end of the barrel region is shown. For the ϕ -strips the ends are beveled instead of the sides.

A normal EM sigma tile is included in the SMD slot to ensure a good EM energy measurement without cross-calibrating the SMD. Offline, the cross-calibration can be used to evaluate the radiation damage and improve the calorimeter energy resolution.

Barrel fiber routing

The WLS fibers are spliced to clear fibers just before leaving the SMD strips. These clear fibers are then routed out of the SMD slot into a 1.5 mm wide fiber-conduit that leads them to the back of the calorimeter. Behind the calorimeter, there are optical-disconnects to provide fiber routing to the photodetectors.

The SMD strips and the fiber routing over their face are shown in Fig. 6-52. Each η fiber turns 90° to come parallel to the top of the slot in order to make the second 90° turn into the 1.5 mm fiber-conduit. The ϕ -fibers must make two 90° turns to come parallel to the slots, followed by a third turn into the fiber-conduit.

These physical SMD strips are one tower long instead of two towers long because we are unable to make reliable fiber connections at right angles to the fiber direction. Such right-angle connections currently lose 50% of the light. Furthermore, their uniformity and radiation-hardness have not been demonstrated. R&D will continue on these right-angle connectors, but we cannot rely on them at this time.




FIG. 6-52. SMD fiber routing in the barrel calorimeter. The 64 SMD strips for a single EM slot are packaged into a rigid box, as shown.

Fiber disconnects

For the SMD, optical disconnects are required because fibers routed through different fiber conduits must be ganged onto the same MCPMT pixels. This would be extremely difficult to do without intermediate connections, especially so when fibers are ganged together. Each 64-channel MCPMT must accommodate 256 fibers which reach the back of the calorimeter in packets of 8 fibers. Special optical cables will be fabricated to transform 32 8-fiber linear bundles at the optical disconnects into 64 4-fiber pixels at the MCPMT. For the SMD upgrade, the optical cables will be replaced with new optical cables which gang two fibers, instead of four, to each MCPMT pixel.

The optical disconnects serve an additional purpose during the fabrication phase for those fibers which connect to a MCPMT on a different wedge. The optical disconnects permit individual wedges to be tested before their final assembly.

Endcap strips

The endcap tower segmentation is shown in Fig. 6-42. The cast-lead endcap EM has bulkheads after every other tower in η . The ϕ -strips are shown in Fig. 6-53; they are one tower long, except at the innermost radii, where they cover only half the tower (0.10 units in η). These strips are not only beveled like the barrel strips, but they are also trapezoidal. The η -strips are shown in Fig. 6-54; they cover half the wedge, or 0.10 units in ϕ , throughout the endcap wedge. As the granularity becomes coarser at higher values of η , the SMD strips become wider (1/8th of the larger tile size), but the electronic ganging is maintained at 0.20 units in η or ϕ . Thus the occupancy levels in the SMD increase by factors of two instead of the factors of four in the corresponding EM tiles.




FIG. 6-53. SMD ϕ segmentation in the endcaps. The physical ϕ -strips are shown in the upper figure; they are one tower long (except for $\eta > 2.6$, where the strips are only 1/2 tower long) and 1/8th tower wide. The electronic channels shown in the lower figure are 0.20 units long and can be upgraded to 0.10 unit.




FIG. 6-54. SMD η segmentation in the endcaps. The physical η -strips are 0.10 units long and 1/8th tower wide. Two channels are ganged to form electronic channels 0.20 units long and can be upgraded to individual readout.

Endcap fiber routing

The fiber routing in the endcap SMD is similar to the routing in the barrel region. Small radii of curvature are imposed on the fibers so they can be brought to the back of the EM through the narrow fiber conduits. The fibers are then routed along the back of the EM modules to the outer radius, where optical disconnects are placed to allow dismounting of the endcap EM.

Electronics channel count

Initially, four physical strips will be ganged together in the barrel region as shown in Fig. 6-50. For the endcaps, two physical strips will be ganged together (four ϕ -strips at the outer radii) as shown in Fig. 6-53 and Fig. 6-54. The total channel count is 47,104, read out by 768 MCPMT's. This system can be upgraded by reducing the ganging by a factor of two and doubling the number of channels and phototubes.

Optical and mechanical R&D

The R&D program for the SMD includes mechanical design, testing of optical components, and radiation-damage studies.

Fiber-routing schemes depend on the final design of the calorimeter. Cost-effective fabrication techniques are still under study. An SMD fabrication station will be built in FY92.

Research to improve the light yield from the SMD strips will continue. The difficulties inherent in the fiber-routing are illustrated by Fig. 6-55, which shows light yield as a function of the bending-radius for a 1 mm fiber. The SMD design requires that bundles of 8 fibers perform 90° turns with radii down to 3 cm.

The delicate position of the SMD near the maximum of the expected radiation damage puts a special emphasis on radiation-damage studies. All materials for the SMD must be tested for resistance to radiation. Fabrication techniques must also be submitted to radiation testing. This will be a major part of the FY92 R&D program.




FIG. 6-55. Transmission of 1 mm diameter fibers as a function of bending radius. The fibers make one full 360° turn at the given radius. The diamonds represent data taken with 1 m of Bicon WLS fiber BCF-91. The squares represent data taken with 2 m of Kururay clear polystyrene fibers. The relative normalization of the two curves is arbitrary.

6.7.5. Photodetectors for SMD

Baseline design

The requirements on the photodetector were described above. The minimum required dynamic range is 4×10^3 , matching the dynamic range of the analog storage chain or the digital design that may be used for the readout. The minimum signal is a single photoelectron.

The best commercially-available device which can meet these goals economically is the multi-channel photomultiplier tube (MCPMT). We have found the 64 channel Philips XP1722 to be the most appropriate for our use. We have chosen a 64-channel tube because our readout geometry is based on groups of 64 channels. Newly-available 96- and 256-channel tubes are poorly matched to the granularity of our system.

The Philips tubes have some limitations:

- 1) The cross-talk between neighbouring channels is 5%. This could give problems for adjacent readout channels from nonadjacent SMD strips.

- 2) The dynamic range of the Philips tube is marginal for our application (10^3). Furthermore, the MCPMT gain sags when there are simultaneous hits in many channels. Correcting the data for this effect could be difficult.
- 3) The rise and fall times of the signals are both greater than 10 ns.

Furthermore, the cost of this tube is relatively high, about \$60 per channel. For these reasons, we consider this tube as a proof-of-principle for our baseline design, but we are pursuing an R&D program to investigate other solutions.

Photodetector R&D

We are investigating three different types of readout for the SMD - MCPMT's, APD's, and silicon target APD-MCPMT's. These different readout techniques correspond to different approaches to development: on one hand, we are trying to improve the existing MCPMT and APD technologies; on the other hand, we are exploring the novel APD-MCPMT technique in search of significantly better performance.

Development of an advanced MCPMT

We plan to explore improved MCPMT designs with Philips and Hamamatsu. In particular, we require improved response to green light (520 nm) and reduced cross-talk ($< 2\%$). As new tubes become available, we will characterize them and evaluate them for our application.

Development of an APD array.

Avalanche photodiodes (APD's) have much higher quantum efficiencies than PMT's, especially for the longer wavelengths necessary for improved resistance to radiation. We expect quantum efficiencies for green light of about 60% for APD's, compared to 10% for green-extended photocathode PMTs. The disadvantages of APD's are that the gain is low (~ 100), the noise is relatively high (at room temperature), and the active surface is usually small. Nonetheless, our recent test-beam results, using individually-potted APD's from EG&G, are encouraging. We would like to develop an APD array to bring the channel cost down to an acceptable level.

We are investigating two technologies for the APD array. A conventional APD array is under study with Hughes Aircraft (Carlsbad, CA), and an original APD array obtained by dividing up the charge collection area of a single large-area APD is being pursued with Advanced Photonix (Los Angeles, CA). These investigations are also a necessary first step toward the development of the APD-MCPMT, described below.

Development of an APD-MCPMT

This device represents a new type of readout technology which we are developing with industry. It combines two familiar techniques to form a MCPMT which should be more economical than tubes made with conventional technologies. The device consists of an array of APD's mounted inside a proximity focussed image tube. A single-channel tube of this type with a large-area APD has already been used to detect photoelectrons. These results shows a clear single-photoelectron peak and a dynamic range of more than 5 orders of magnitude. We believe that this technology could produce a high performance device for half the cost of a conventional MCPMT.

There are two areas which need to be developed for this device. First, an 8×8 APD array must be developed. This line of development, carried out with Hughes Aircraft and Advanced Photonix, has been described above. Then, the array must be potted in a suitable tube equipped with a photocathode. We have discussed with Litton Electron Devices (Tempe, AZ) the problems associated with manufacturing such tubes. Based on the experience gained with the single-channel device, they feel that this new device should not present any major technological difficulties.

Development of VLPC array

The visible light photon counter (VLPC) could offer a number of advantages. It has high gain (10^4) and immunity to magnetic fields, and should have high quantum efficiency and no cross talk. Preliminary measurements indicate a dynamic range of at least 3000. In an ongoing effort for scintillating fiber tracking, Rockwell International is developing large arrays of VLPCs and engineering the necessary cryogenics (Chapter 4).

6.7.6. Calibration of the shower maximum detector

The calibration program for the SMD includes source testing during fabrication, quality control during installation in the EM, and calibration *in situ*. We must verify that every strip responds correctly before they are sealed into the calorimeter. Gamma rays from sources in our calibration tubes have sufficient range to satisfy our calibration needs. The general calibration scheme is discussed next.

6.8. Calibration

6.8.1. Introduction

The tile/fiber choice of calorimetry necessitates a precise calibration scheme since, unlike ionization detectors, it is inherently neither uniform nor stable. Calibration systems are needed to:

- Perform quality control during assembly. The uniformity within tiles, the design of the mask to establish uniformity among tiles within each tower, and the tower-to-tower response variations need to be determined.
- Monitor the performance of PMT's and front-end electronics, both in gain and linearity.
- Transfer the test beam calibration of a subset of modules to the rest of the modules at startup.
- Maintain calibrations from a test beam or collider data. Both absolute and tower-to-tower gains, including the effects of radiation on the optical system, need to be tracked in time.

Two types of systems will be used in conjunction with test beam and collider data: source delivery systems and optical flashers.

The source delivery system can expose each individual tile in the calorimeter to a radioactive source. The induced current measures the conversion of ionization to light, the optical transmission, and the PMT gain. This system can be used both in quality control at assembly and in determining the detailed profile of light loss due to radiation damage. The detailed measurements can be used to transfer the behavior of those modules calibrated in a test beam to those modules not tested in a beam. The stability and reproducibility of the source system allows the maintenance through time of an existing calibration.

A light flasher system is used to complement the source system; the source creates a current and does not test the pulse-specific aspects of the readout. The flasher system will be used to check PMT and front-end electronics gain and linearity, and may also be used to check a subset of the optical system. On-board electronics calibrations should be sufficient to factor out front-end electronics performance (see chapter 8).

A subset of calorimeter modules with calibration systems will be put through an extensive and systematic program of calibration and response studies in a test beam (see chapter 14). The study of the relationship between beam response and detailed source measurements in several modules will allow prescriptions to be developed in order to extrapolate to modules exposed only to sources.

The ultimate absolute and relative tower-to-tower calibrations will come from collider data. In a month of running at design luminosity, a sufficient number of electrons from W and Z decays will be recorded to absolutely calibrate individual EM tower gains to $\pm 0.2\%$.

6.8.2. Source calibrations

Each row (constant ϕ) of scintillator tiles in each depth layer of the central calorimeter is faced by a source tube. In the EM sections these tubes are cast into the lead absorber, while in hadron sections the tubes are located in grooves milled into the iron. Radioactive sources mounted on a wire are moved through the tubes while the current outputs of the PMTs are monitored. The same source can be used for all tiles in a wedge. During assembly, and *in situ* with the endcaps pulled back, all of the tubes are accessible to source drivers to enable detailed response mapping for each tile. When the endcaps are in place, only a fraction of the tubes are accessible. In particular tubes facing two of the tiles in each depth readout (EM1, EM2, HAC1, HAC2) are available at all times for stability monitoring and calibration maintenance. These two functions—tile mapping and stability monitoring—may require different sources.

For mapping, the low energy photons from ^{57}Co may be most appropriate. In the quality control phase, each tile is tested and the tile-to-tile response within each tower is equalized by masking to $\pm 2\%$ in the EM and $\pm 5\%$ in the HAD sections. This detailed information is retained for use in transferring the test beam calibration. The same detailed measurements will be repeated approximately once per year to determine the depth profiles of light loss due to radiation damage. For currently available plastic, one year of running at design luminosity should produce a light loss of $\sim 20\%$ in the worst case, EM shower-maximum at $\eta = 3$. (The loss for $|\eta| < 2.5$ will be $< 5.5\%$ during this period). The detailed map can be used to optimize gain coefficients, redesign and replace longitudinal masks, and determine which, if any, tiles need to be replaced. Once the damage profiles are well measured, source results from the tubes which remain accessible when the endcaps are closed will be more readily interpretable.

The accessible tubes will be used to maintain calibration over time; ^{137}Cs , with its long half life, is more appropriate for this function. Calculations and measurements show a decrease in response of $\sim 2\%$ as the source moves 0.012 in. away from the tile. Location of the source to this accuracy could be achieved by the smaller of the source tubes used on the CDF plug upgrade and SDC EM prototypes or by a bead or collar on the end of the source carrier in the tubes [40]. We expect, for our PMT gain, a current of a few nanoamps per milliCurie of ^{137}Cs source activity. We would, therefore, need to measure this current to an accuracy of a few $\times 10$ picoamps. The ^{137}Cs source system used for the CDF central EM calorimeter has been demonstrated to be reproducible to $\sim \pm 0.7\%$ [43]. A systematic study of effects of magnetic field will be needed in order to disentangle scintillator response changes from source effects, and to correctly interpret the results from tile mappings performed with the endcap open and the field off.

6.8.3. Optical flasher calibration

We are developing a laser light source and variable optical attenuator system for illumination of the PMT's and/or calorimeter tiles. The objective of such a system is to monitor the the linearity of response and channel gains over extended periods. This system is at an early stage of development but basic requirements and a number of implementation issues have been identified. First, we must understand the accuracy requirements on the linearity and, if the tiles are to be illuminated, which subset of tiles need to be monitored. Monte Carlo studies are underway to determine the effects of illuminating only a fraction of the tiles. The gain monitoring function requires that the intensity of the laser source be well determined. This can be achieved by the use of a photodiode system.

Work will begin shortly on the configuration of the fiber system to deliver light to the PMT's and, if necessary, the tiles. The issues here are the choice of fiber, routing to minimize the size of calorimeter penetrations and fiber lengths. We shall also investigate the optimal manner and location for coupling the fiber ends to the tiles e.g. in relation to the orientation of the readout fiber in the sigma tile design. As discussed above, other elements of the calibration system will provide channel-to-channel relative gains, and there is therefore no severe requirement to equalize light output. The only requirement is that the light from each fiber covers the necessary dynamic range. The fiber length equalization will need to be regular enough for the 16 ns gating, and possibly to allow the laser to provide a timing standard.

A laser has been procured for this work and the initial aim is to provide a calibration laser system for the 1993 calorimeter prototype module. We shall also study the possible use of the laser system to

provide timing synchronization signals to all parts of the calorimeter; the issues here are laser pulse shape variations and time stability.

6.8.4. *In situ* calibration

Relative (tower-to-tower) and absolute calibration can be obtained from collider data.

The ratio of energy measured in the calorimeter to momentum measured by the tracker (E/p) of electrons from $W \rightarrow e\nu$ and $Z \rightarrow ee$ events will be used to establish absolute calibration of the EM compartment. An absolute individual tower calibration to the level of $\pm 0.2\%$ can be achieved in approximately one month for towers with $\eta < 2.5$ [9]. Towers at larger $|\eta|$ benefit less from magnetic tracking. They will need to be calibrated with $Z \rightarrow ee$ events using the invariant mass constraint, which will take longer.

For hadron compartments, absolute normalization can be established, especially at low energy, using E/p for isolated tracks (π^\pm). The background from overlapping π^0 's must be subtracted [42].

The absolute EM compartment calibration can be transferred to the hadronic compartment using γ -jet E_t balance. The opposing jet can, therefore, be used to transfer absolute calibration to the hadronic towers [42].

The entire dijet sample can be used to obtain the relative jet response as function of η , using E_t balance. The upper range of energies accessible using this method is limited by kinematics in the forward direction, but the calibration is more important at lower energies where nonlinearities appear. Assuming no gross nonlinearities of response to energy, hadronic calibration can be done to the 3% level on the interval of a few days.

The W and Z electrons will span a reasonable range of p_t , allowing study of the EM1/EM2 balance and linearity.

Using some of the techniques of this *in situ* approach and a relatively small data sample of 1500 W electrons, CDF achieved an absolute calibration of the electron energy scale to $\pm 0.24\%$. A relative calibration of 480 EM towers to $\pm 1.7\%$ was achieved using 17000 inclusive electrons.

6.9. Engineering and R&D plan for the central calorimeter

6.9.1. Introduction

The R&D program for FY92 and beyond has as its aim the establishment of a consensus design choice, engineering work on a preproduction prototype for the barrel region, and continued work on the basic properties of the detection medium. The goal is to have the prototype constructed, fully outfitted and in a test beam during 1993.

The program of radiation hardness studies and development of improved radhard-optical systems discussed in Section 6.5 could have important impact on the design of the endcap/plug calorimeter. The SMD related R&D discussed in Section 6.7 is equally important; a cost effective optical transducer needs to be developed since no completely acceptable candidate now exists.

The baseline design favors technically lower risk design options with understood costs. The optimization program over the next few months will be to review this design for physics performance in light of manufacturability and cost (value engineering). On the one hand, more cost effective design options will be found and validated. On the other hand, a more detailed cost development may reveal added cost elements not presently included.

The engineering plan describes a set of tasks necessary to refine and complete the baseline design described in this report. Each of these tasks will require decisions based on analysis of performance, cost and manufacturability. Performance issues are being addressed by particle beam tests and computer simulations. Cost information is being gathered from vendors having experience with related projects. Manufacturing costs are derived from building prototype components in the laboratories or in industry and

scaling the actual costs appropriately. Instrumented prototype detector modules will be built in FY92 and tested to verify the mechanical design, to gain experience in assembly and to study performance. During FY93, prototypes will be built and manufacturing plans will be made for a complete barrel wedge assembly.

6.9.2. Engineering tasks

Absorber structure

The basic method for producing the hadron absorber structure outlined in this report is to stack and weld machined plates. Options still exist, and are being pursued in parallel with detailing the design of the baseline proposal. One example is the technique of stacking and welding with the use of spacer plates; this may prove to be better matched to the available manufacturing resources and result in lower cost.

An extensive engineering R&D program resulting in a full scale cast EM absorber section is currently underway. It includes tests with various alloys using both small and larger castings which will explore techniques for endplate and bulkhead connections. Additional studies are being planned in order to understand stress and deflections, mold design, fabrication, and support frames.

Details of the barrel absorber configuration such as source tube routing and fiber routing grooves need to be specified. The placement and sizing of all bolts and welds and the connection of modules need to be determined in the context of support structure distortion and impact on hermiticity and tile placement. Structural analysis of the absorber structure needs to be done taking into account the gravitational, magnetic, transport and assembly forces to which the modules will be exposed. The absorber structure must be compatible with other detector subsystems, and details of these interfaces need to be determined. Examples include location of the solenoid chimney, and solenoid and tracking system support brackets.

The endcap absorber design details must accommodate proposed tile replacement scenarios. The endcap/end plug interface needs to be carefully specified as do other module interfaces. Designs for the connections between the barrel halves, and between the barrel and end caps need to take into account the existing forces and interactions with the support structure. Details related to handling all modules, including lifting points and access requirements need to be determined.

Tile/fiber readout

The basic issues of the tile/fiber readout system have been studied in various R&D programs and exercised in prototypes. One of the exceptions is the choice of a radiation-hard plastic for the forward sections of the central calorimeter. R&D will continue on all fronts with emphasis on as yet undetermined aspects of the design and optimization of manufacturing techniques and minimization of cost. Aspects of tile geometry such as the fiber groove cross section and fiber routing need to be finalized for maximum light yield and uniformity. Satisfactory methods for producing and installing the tile assemblies have been exercised in the various test modules but require optimization in terms of manufacturability and cost. The ongoing R&D program will determine the diameter and type of waveshifter and clear fibers that will be used. The process of splicing clear and waveshifter fibers exists, but needs further development with regard to volume manufacturing requirements. Upgrade and replacement scenarios require the use of optical disconnects which need further development. Fiber routing and attachment along the absorber will require further study with the goal of facilitating assembly and at the same time minimizing cracks.

Photomultiplier tubes

Candidates for the photomultiplier tubes need to be evaluated against the requirements. The layout of the PMT's and the connection of the assemblies to the absorber need to be determined with consideration of replacibility, access and upgrade scenarios. The PMT assemblies require design modifications to be able to take advantage of cost effective manufacturing options such as the use of injection molded parts. The specification of the PMT gain monitoring system needs refinement to help improve input to the cost estimate.

Electronics crates

The number of electronics crates and the method of attachment to the absorber need to be determined. Cable routing paths and crate placement will be designed with consideration of required PMT access. All crate utilities must be specified and routed along the absorber. The layout needs to be designed preserving the modularity of the absorber structure. If crates on adjacent modules share utilities, the connections for the utilities must be designed to separate easily along module boundaries.

Radioactive source calibration

The technique of radioactive source calibration is currently the subject of intense study. There are several engineering aspects which need attention: the method of attaching the source drivers to the calorimeter while allowing access to the PMT's, source tube manifold and connections for the straight tubes accessed only during shutdowns, and the mechanics and control system for the *in situ* monitoring subsystem.

Support structure

The support system must be designed to hold up the calorimeter under all conceived loading conditions, specifically accounting for distortions caused by the muon steel deflection and tolerance accumulation. The design of the interface between the supports and the absorber should account for utility penetrations and required access. A modular structure attached to the muon steel beneath the calorimeter must be designed for the calorimeter to roll on during its positioning in the toroid. This activity should be done in conjunction with the muon group. The final design of the support system should include all necessary alignment degrees of freedom (roll, pitch, yaw) as well as horizontal translation, parallel and perpendicular to the beam, and vertical adjustment. The systems necessary to control the motions of the calorimeter need to be designed. These include the system to drive the calorimeter into the muon toroid and the hydraulic alignment control system.

Tooling and equipment

Once the basic component design becomes more detailed, the engineering effort will begin to focus more closely on the necessary tooling and equipment required for assembly of the calorimeter and components. General categories of tooling include fabrication and inspection tooling to be used for the tile/fiber, support structure and absorber component assemblies. Lifting fixtures are needed for absorber modules and support structure components. Equipment is needed for splicing and polishing fibers, facing PMT cookies and electrical testing of PMT's. Tooling and fixturing are required for tile insertion, fiber routing, PMT housing assembly and inspection. Erection tooling includes devices for assembly of the support structure, wedge assembly, the lowering platform and a fixture to insert/extract the plug from the end cap.

Facilities

Facilities for fabrication and assembly of the calorimeter components, steel and lead absorber, tile/fiber assemblies and module assembly site(s) need to be assigned and specified. The surface facility at the SSC site must be designed and the lowering and installation procedures in the collision hall outlined.

Cost estimate

Implicit in the above list/discussion is that a more robust cost estimate will emerge as the design becomes more detailed. In the near future, attention will shift to optimization of existing techniques with emphasis on cost reduction, manufacturability and reliability.

6.9.3. Safety analysis

A hazard analysis of the central calorimeter system is under way. Flammability of the plastics, oxygen deficiency associated with a possible nitrogen purge system, radiation from the source calibration system and beam activated absorber material are the primary hazards. These hazards will be studied and minimized in the design, operation and maintenance of the calorimeter system. Manufacturing operations will be carried out so as to ensure the safety and health of the workers and also to ensure minimum impact of fabrication operations and materials on the environment.

6.9.4. Cost estimate development

The principal cost drivers in the central calorimeter are the absorber assemblies, the tile/fiber optical system, the photomultiplier assemblies and the source calibration system.

The process of casting lead radiator arrays has been studied extensively by ANL and the Westinghouse Science and Technology Center. Several model assemblies have been built under the joint direction of ANL and WSTC. As a result, cost information for this process is fairly well understood.

The steel absorber assembly has also been costed using the method of stacked, machined plates, plug welded together. Vendor estimates are based on detailed drawings. The experience with the fabrication of similar large steel structures for CDF has added confidence to cost estimates. Close contact with vendors will continue to further optimize the absorber fabrication method.

Alternate assembly techniques for the absorber are being explored in an effort to reduce cost. One method builds up steel absorber using stacked, punched laminations. This method is used commercially to fabricate power transformers and has been used extensively in magnets for accelerators including the SSC. The difficulty with this scheme is finding a method of joining the laminations into a strong rigid block suitable for finish machining. The use of epoxy was found to be inadequate. Fusing of the plates in a brazing oven in a hydrogen atmosphere shows promise. Another possibility is the use of thick laminations (approximately 2 inches thick) in which the tile slots are cut with an abrasive water jet rather than a punch. A laminated absorber could result in significant cost savings.

Candidate methods for tile fabrication are machining and injection molding. Machining from standard scintillator plastic sheet using CNC machines seems to be the most straight-forward and cost effective way to make the bulk of the tiles, and this method was assumed in formulating the cost estimate. *There remains* some concern, however, about stresses and possible long-term crazing induced by machining. The CNC machining method has the overwhelming advantage of being able to machine with ease the large variety of tile shapes required. Molding may be used for the special radiation hard tiles because of the composition of the plastic required.

Because of the large number of phototubes (approximately 20,000), their cost impact is sensitive to marginal changes in the unit cost. Fortunately recent phototube assembly cost experience with the large detectors provides a reliable guide for the PMT costs.

The cost estimate for providing a complete source calibration of all tiles during calorimeter operation in the detector was recently completed. It was found that the cost of machining grooves with compound curves in the steel to accept the source guide tubes with the endcap calorimeter in place has an unacceptably high impact on the cost of the absorber. The scheme proposed here features on-line source access to only a few tile layers in each tower. Detailed measurements as needed due to radiation damage to tiles will be made when the end caps are drawn back, allowing access of source tubes for each tile. In this scheme most source tube grooves are straight in the absorber wedges.

6.9.5. Prototypes

Verification of design properties, cost and performance depends on the construction of prototypes of components and complete subassemblies (such as wedges). The scale of these prototypes must be adequate to address the crucial issues. The SDC collaboration has already constructed and tested a series of absorber structure models. Instrumented cast lead EM test modules have been built and tested in a particle beam. Steel absorber structural test models have also been fabricated using different construction methods. A steel absorber module with two instrumented towers has been placed into a test beam. The design will become firm, and a full wedge prototype will be constructed and tested.

Prototype hadron wedge

A portion of an SDC barrel hadron half-wedge has been fabricated using the method of stacked and welded plates with spacers. The module is 2 towers in ϕ and is truncated to 4 towers in η starting from the 90° end. The plates were stacked starting with a 4" thick "bottom" plate which would become the top or backplate of the module. The plates were then successively stacked and plug welded. There was no significant warping due to the welding. A slight reduction in plate thickness (10-15 mils) due to a light grinding pass used to clean the plates resulted in the module being 1" shorter than designed, an error which can easily be compensated. Systematic variations in the plate thicknesses were controlled by grinding the plates to compensate. The two faces with tile slots were machined and the grooves for the readout fibers were cut without any particular problems. The block was then deburred and cleaned. Each tile slot was deburred and swabbed to remove the cutting oil and chips. The source tubes were prepared by cutting to length and shaping the ends for the flex tube fittings. Many of the time consuming steps used in the preparation and setup of the source tubes and driver could be improved by using plastic tubing throughout. This option is being investigated. Two towers each of HAC1 and HAC2 tiles were installed. Installation of the tiles was quite straightforward. Each tile was held against the source tube by a simple aluminum "spring" which worked very well. Most of the time was spent routing and dressing the fibers, as anticipated. The consensus reached by those involved in installing the tiles was that new techniques aimed at routing the fibers in a single layer should be investigated.

The test wedge experience was a valuable guide in extrapolating to a full-scale production module. The plate and spacer stack assembly procedure for the steel block, though it yielded good results, is probably too time consuming (i.e. expensive) and involves the handling and tracking of many, many separate and unique parts. Alternate methods for producing the absorber, such as the current baseline proposal of stacking machined plates, promises to be simpler and less expensive. The possibility of stacking thick laminations is still a viable alternative at this time and reduces the number of parts down to only 8 per wedge.

The handling of this relatively small block (5 tons) was fairly simple and made even more simple since it did not have hermiticity requirements (edge geometry was not important). Handling a full sized half-wedge will be a more difficult problem as one needs access to all sides except the two ends. In addition it must be able to rotate easily around the longitudinal axis.

Fiber routing was as difficult as we imagined it might be, and more detailed work on routing and packing the fibers in the groove should be done even before we embark on a pre-production prototype.

Potting the fibers and facing off the cookie went well despite the obvious need to gear up to production techniques. Work still needs to be done on fixturing and techniques which will minimize the stress on the fiber bundles and fiber/cookie interface.

The PMT assembly was straightforward but there are several changes that would be necessary (mostly cost related) for the full scale device.

Cast lead EM absorber

The fabrication method for the lead EM absorber material is based on large-size alloyed lead castings poured around structural support skeletons of welded stainless steel plates. In particular, stainless steel bulkheads are positioned between pairs of scintillator towers with the cast lead plates constituting the parallel planes of passive absorber interleaved with scintillator plates to form the calorimeter structure. A modular group of (barrel or endcap) towers constitutes a wedge-shaped calorimeter module that is the basic fabricated subsection of the calorimeter. There are 64 endcap and 128 barrel castings in the complete central calorimeter.

To verify the properties of a cast lead EM absorber, two 10-tower correct-geometry test castings were produced in 1991, filled with scintillator plates and tested in a tagged beam of variable-energy electrons at Fermilab (Section 6.11.2). The cast lead absorber met all performance criteria. Furthermore, the actual costs of fabrication using the cast lead method compared favorably with a stacked plate method used earlier for the CDF calorimeter and produced a module with superior overall dimensional and hermetic properties [44].

Because pure lead is a material with poor mechanical properties (it especially suffers from creep), extensive investigation was completed on finite-element analysis of stress in the composite lead/iron wedges, and laboratory tests were made at ANL of the creep rate at elevated temperatures. These tests indicate that pure lead can be used in the present design for the EM calorimeter, but the margin against excessive creep is narrow. As a result, we will produce a new cast lead 10-tower test module in FY92 with thinner absorber plates (4 mm) out of calcium/tin alloyed lead that has been shown to exhibit much greater resistance to creep. Such alloys are routinely used in commercial battery plates and we have already obtained a sample of a suitable alloy. It remains to be demonstrated that the same casting technique followed with the test module using machined aluminum inserts will be as successful with the alloy as with the pure lead material.

A design for a full-size barrel EM casting will be completed in FY92 in expectation of casting a full-size prototype in FY93. The scale-up of the 10-tower test module to a 56-tower full size prototype should be straightforward.

6.9.6. Manufacturing plan

The fabrication of the central calorimeter will be similar to other projects successfully completed in the national laboratories. Components that can be made better and cheaper in industry will be made there. Operations requiring special skills not readily available in industry will be performed at the laboratories. The steel and lead absorber fabrication, tile cutting and wrapping, fiber termination and assembly tooling fabrication will be done in industry. The actual assembly of the tiles into the calorimeter, routing of the fibers, connection to the photomultiplier tubes, masking and component performance testing will be done in the national laboratories.

Quality Assurance/Quality Control will be an integral part of the manufacturing plan both for in-house assembly operations and for fabrications of components in industry. Development of the QA Plan and QC procedures will be done in parallel with the detailed engineering design of the calorimeter elements and subsystems.

6.10. R&D and engineering for the forward calorimeter

In the section below we briefly describe the R&D and engineering plan for the two options for the forward calorimeter: high-pressure gas and liquid scintillator.

6.10.1. High-pressure gas ionization option

This technique has unity gain and is expected to be extremely radiation hard. Detailed tests are planned to evaluate stability under intense radiation. The safety problem can be solved by constructing the calorimeter out of small modules. This approach has been tested and proven to be very safe as discussed later in this section.

Radiation damage tests with a high-pressure vessel

Argon gas at 100 atm inside a test vessel has been exposed to about 16 Mrad of radiation from a ^{60}Co γ source. The signal from an ^{241}Am α source inside the vessel was measured before and after irradiation. Only a 6% reduction in the signal size was observed after the exposure [46]. After the exposed gas was purged and the vessel refilled with new argon gas, the α source signal returned to its value before irradiation. A similar test was done with the test vessel filled with 100 atm argon mixed with 1% methane. No significant change in the α source signal after irradiation was observed [46].

Results from beam tests of a prototype EM calorimeter

A prototype, parallel plate electromagnetic high pressure calorimeter has been built and tested. This detector consists of ten high-pressure vessels, each of which consists of two steel plates bolted together. The thickness of each plate is 1.25 inches and its diameter is 12 inches. A 10 inch diameter, 5 mm deep groove is milled in one of the two plates. A 1 mm thick copper clad G-10 readout board is placed in the groove and supported with ceramic spacers. High pressure gas occupies a 2 mm gap on both sides of the G-10 board. This calorimeter was tested in a test beam at Fermilab. Most of the data were taken with the calorimeter filled with 95% argon and 5% methane at 100 atm. Other gas mixtures were also tested. Some preliminary results [47] are:

- a) The signal from the EM showers saturates at the very comfortable voltage of 1.4 KV (or 0.7 KV per mm).
- b) The calorimeter is linear in the energy range tested (20 to 125 GeV).
- c) The measured energy resolution agrees well with an EGS4 Monte Carlo prediction.
- d) The calorimeter signal is proportional to the gas pressure in the pressure range tested (20 to 100 atm).
- e) The inverse of the drift velocity in 100 atm argon plus 5% methane is 20 ns per mm. This gives a total signal collection time of 40 ns for a two mm gas gap.

In summary, although the geometry planned for the forward calorimeter involves tubes rather than parallel plates, it has been shown that a high-pressure gas calorimeter can be operated, is radiation hard at least up to 16 Mrad, and is fast enough for our application.

Proposed calorimeter design and R&D for FY92

In FY92 a combined EM-hadronic prototype calorimeter will be designed, constructed, and tested. This design is based on tubes oriented at a small angle relative to the shower axis. GEANT Monte Carlo studies [48] show that the stochastic and constant terms of the hadronic energy resolution will be 70% and 4%, respectively, well within SDC specifications. The prototype calorimeter is square in shape and consists of sixteen $10 \times 10 \text{ cm}^2$ modules. Each module is made of 81 tubes closed at one end and welded to a one inch thick steel manifold at the other end. The absorber consists of steel laminations, similar to the ones used in magnet yokes, with a matrix of punched holes for the tubes to pass through. The tubes are 3 m in length. Charge collection rods run along the center of the tubes. The gap between the rod surface and the inner tube wall is 2 mm. The final inner tube diameter and the rod thickness will be chosen, after extensive Monte Carlo simulations, to optimize the energy resolution. At this point our best estimate for these parameters is 7 mm and 3 mm respectively. Several ceramic spacers are placed along the rods to minimize sagging, and possibly to reduce the electromagnetic resolution constant term. A ceramic spacer is placed at

one end of each rod to avoid electrical breakdown between the rod at positive high voltage and the tube, which is being held at ground potential. The rods are all connected electrically with a wire or a copper clad G-10 board at the other end. If the electrical noise proves to be too large because of excessive source capacitance, a smaller number of tubes will be connected together. A steel plate of a few mm transverse wall thickness is welded to the manifold. This plate, being on the side away from the incident particles, can be as thick as needed. It carries the gas and the high voltage/signal feedthroughs and, if needed, the gas relief valve. The module will be designed to operate at 100 atm. In a hydraulic test performed on the manifold and the end plates welded together, the welding failed when the water pressure reached 640 atm; this is well above the design operating pressure of 100 atm. A forward calorimeter can be built from about 1600 such modules as shown in Fig. 6-56.




FIG. 6-56. The forward calorimeter as realized in the high-pressure option.

To summarize, the 1992 R&D effort of the group working on high-pressure gas calorimetry will focus on designing, building and testing a prototype, spaghetti type calorimeter, on developing a low noise and fast preamplifier that best matches the characteristics of the proposed modules, on conducting further radiation damage tests, up to 100 Mrad, on the proposed coaxial geometry under high voltage, and on producing an engineering design of a full scale high-pressure forward calorimeter.

6.10.2. Liquid scintillator forward calorimeter option

The liquid scintillator/capillary concept aims at achieving a speed well matched to the central calorimeter systems, and at the same time realizing a very radiation hard technology. The speed of the scintillator system, the concomitant lack of pileup problems, and the possibility of utilizing the same electronics throughout the calorimeter system, are strong points. However, it has to be shown that the proposed system meets the requirements on radiation hardness.

The conceptual design for the liquid scintillating option consists of a steel absorber matrix interspersed with 5 mm diameter quartz tubes in towers as shown in Fig. 6-57. Liquid scintillator flows through the quartz tubes, at a rate such that the necessary level of survivability is attained. Central to this concept is the idea that each azimuthal quadrant of the forward calorimeter system is contained in a single vessel. This design has been chosen to minimize the extent of liquid scintillator plumbing, and also to allow one to achieve the necessary liquid flow by means of differential pressure within each vessel. The scintillating tubes are read out by clear fibers coupled to the upstream ends. The clear fibers are brought out through

the free standing liquid surface to the photodetectors. It is a design aim that the liquid containment vessels have no feedthroughs below the free standing liquid surface, in order to minimize the problem of leaks. Implementing the readout with clear fibers allows a very flexible approach to the spatial division of the device. However, it does raise the problem of flowing the scintillating liquid into the readout ends of the tubes. We are working on the mechanical design issues. We plan to start building a hadronic module, incorporating all the necessary mechanical features, by the end of 1992.




FIG. 6-57. Conceptual Sketch of Pointing Tower for Liquid Scintillator FCal (outer fibers terminate at the angled tower surface, to maintain constant sampling fraction).

The argument that radiation hardness should be a natural feature of liquid scintillator capillary calorimeters is based on two simple ideas. Liquids are, by their very nature, more radiation hard than plastic materials. Moreover, liquids can also be circulated, or completely replaced, very easily. Thus the scintillation medium can effectively be made radiation hard. The impact of radiation damage arises from transmitting light through a solid transparent medium. However, in the liquid capillary concept, the light from the liquid is trapped by total internal reflection at the capillary walls. Since the evanescent wave only penetrates of order a wave length into the capillary wall, the total light path length in the solid is of order a centimeter, rather than typically a meter in a solid optical calorimeter. Hence, for the capillaries, a radiation damage induced reduction in the attenuation length of the solid medium is much less important than in tile/fiber calorimeters.

There are essentially two approaches to the capillary question. One can use scintillators with refractive indices of 1.5, which leads to the requirement of low refractive index $n = 1.4$ tubes. A possible combination is NE224 in a Halar or Viton (fluoropolymer) tube. This has been shown to be radiation hard at the level of 10's of Mrads [49]. However, this system has to be investigated at the level of 100's of Mrads for use in the forward calorimeter.

Radiolysis is also an issue to bear in mind. The irradiation of scintillators may well produce radiolysis products which coat the inside of the tubes, and thus dramatically reduce the efficiency of internal reflection. It could also produce corrosive substances which attack the tube walls. Outgassing of the solvent is also a possible problem. An alternative approach to the scintillator tube combination is to use quartz $n = 1.46$ tubes containing an isopropyl-biphenyl based scintillator $n = 1.58$. This combination is interesting in that it has the same refractive index combinations as scintillating fibers composed of polystyrene cones with PMMA cladding. Both quartz and isopropyl-biphenyl are known to be radiation hard materials. Isopropyl-biphenyl is known to be radiolysis resistant, having been investigated as a candidate reactor cooling liquid. NE Technologies produces a green emitting radiation hard scintillator based on isopropyl-biphenyl.

In the coming months the liquid-spaghetti group will pursue a program aimed at demonstrating that liquid scintillating capillaries can be used in the calorimeter, and that sufficiently long attenuation lengths

can be maintained in a high radiation environment. At least one EM prototype will be manufactured, and tested at CERN. In addition, schemes for using clear plastic readout fibers without interfering with the liquid scintillator circulation will be investigated, and the detailed mechanical design of a realistic system will be pursued.

6.10.3. Milestones for 1992 and 1993; the next decision level

Both technologies will be supported to: a) build and beam test a credible hadronic prototype; b) pursue decisive radiation damage tests; and c) produce a sufficiently detailed engineering design so that the compatibility with the requirements can be judged, and a serious cost comparison can be made.

It is anticipated that reduction to one technology can be accomplished in late 1993.

6.11. Test beam results (FNAL T841)

The development of the designs for the SDC calorimeters will continue to require beam tests. The overall SDC test beam plan is discussed in Chapter 14. Here we review two test beam efforts that have been quite influential in design issues discussed above.

6.11.1. Beam test of reconfigurable-stack calorimeter

A "Reconfigurable-Stack (Hanging File) Calorimeter," a large test beam calorimeter which contains lead, iron, and aluminum plates, has been built using scintillating tiles as the active media. It can be reconfigured to vary sampling fractions, material ratios, and low- Z cladding layers. The primary objective of the hanging file calorimeter is to identify practical combinations of materials, thicknesses, and sampling fractions that yield acceptable resolution and compensation for the SDC hadron calorimeter. Beyond this main objective, another goal is to provide data, with a variety of absorbers in a consistent set of experimental conditions, to tune up the various hadronic cascade Monte Carlos. An example is given in Fig. 6-58.

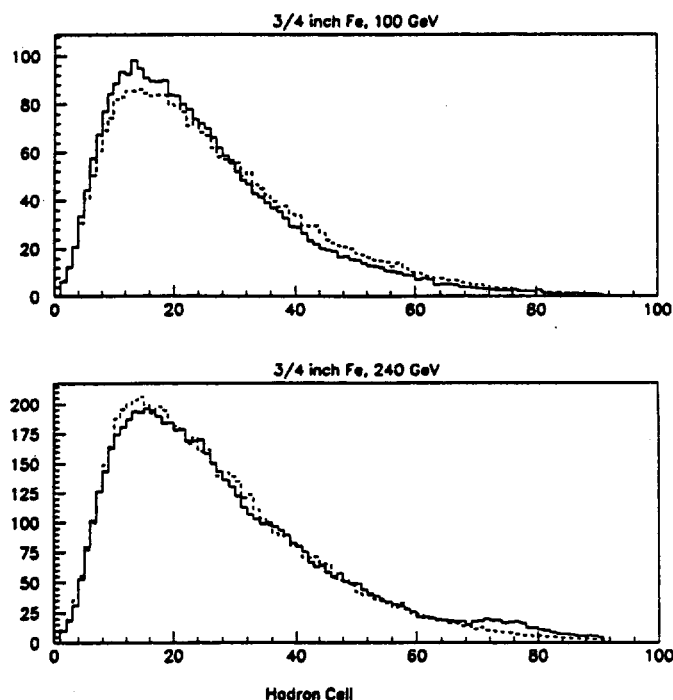


FIG. 6-58. Comparison between the longitudinal energy deposition predicted by CALOR89 (solid histogram) and hanging file test data (dashed histogram). Test configuration is a 0.75 inch iron absorber followed by a 3 mm scintillator, repeated 92 times. The top histogram is for 100 GeV π^- and the bottom is for 240 GeV. The vertical scale is arbitrary.

Description of reconfigurable-stack calorimeter

The calorimeter stack is formed by a series of $1\text{ m} \times 1\text{ m}$ plates which can be any combination of scintillator tile, lead, iron, or aluminum. Since more than 95% transverse containment of hadron showers is achieved in a $60\text{ cm} \times 60\text{ cm}$ area, the transverse dimension of the stack provides a comfortable safety margin. With absorbers as thin as $1/8$ inch to be tested, up to 600 plates can be stacked to form a $\sim 10\lambda$ deep calorimeter. For the cladding tests, 200 $1/8$ inch and 100 $1/16$ inch aluminum plates were used. Two identical box-beam structures similar to a hanging folder file, were built to hold the plates. One structure was used to hold the configuration for data taking, and the other to allow the plates to be inserted and removed from the beam enclosure for restacking and for temporary storage.

The readout plates are 3.0 mm thick tile/fiber assemblies. Each plate is read out by 20 1-mm fibers, running the full length of the tile to ensure uniformity. Each tile is individually light-tight and is read out by its own phototube. Radioactive source tests showed that the transverse uniformity of $\sigma \sim 2$ to 5% was achieved in the central region ($60\text{ cm} \times 60\text{ cm}$) of each tile.

Muons were used to obtain the first-order calibration constants for each plate. An absolute light yield of ~ 1.3 p.e./MIP per plate was observed. The energy resolution and shower profile of the electron data are in excellent agreement with predictions of the EGS Monte Carlo. The resolution of the pion data agrees well with previous experiments where a comparable configuration was used.

Preliminary results for SDC calorimeter design

During the 1991 fixed target running period at FNAL, data were taken with 40 different configurations of materials, thicknesses, sampling fractions, and low- Z cladding layers over the energy range of 5 to 240 GeV. The linearity, resolution and e/π response are shown in Fig. 6-59 for two configurations which are relevant to the choice of absorber in the SDC hadronic section. To obtain $e/\pi \sim 1$, a different relative weighting of signals from the EM and hadron compartments was used for each configuration. Fig. 6-59(a) shows the linearity and Fig. 6-59(b) shows the pion resolution after the proper weighting. In Fig. 6-59(c), the measured e/π response is shown as a function of energy. The data indicate that the lead and iron hadron calorimeters have comparable performance.

Table 6-11
Tested configurations for SDC calorimeter.

Configuration	EM (# layer)	Hadron (# layer)	Depth (λ)
I	$1/8$ in Pb (40)	1 in Pb, (52)	9.0
II	$1/8$ in Pb (40)	1 in Fe, (55)	9.5

6.11.2. The MP program

The Fermilab MP test beam has been used to study two electromagnetic calorimeter tile/fiber prototype modules, the scintillating fiber preshower detector prototype, and the shower maximum tile/fiber prototype. In the case of the EM calorimeter, the goals of the test beam work were to demonstrate the viability of the calorimeter, verify that the resolution and linearity were close to design specifications, and study some detailed questions such as calorimeter response variations near module boundaries, over the area of a tower, and for different angles of incidence of the showering particle. In the case of the preshower and shower maximum detectors, the goals were to demonstrate the ability to localize showers with millimeter resolution, to distinguish photons from π^0 s and to separate charged pions from electrons. A further goal of the preshower (massless gap) is to provide an energy correction to the EM calorimeter for energy lost in the magnet coil.




FIG. 6-59. Measured (a) linearity, (b) pion resolution and (c) e/π response as a function of energy from the hanging file calorimeter configurations shown in Table 6-11

The EM calorimeter resolution was found to be consistent with the prototype design target of $18\%/\sqrt{E}$, using electrons with energy 15–35 GeV (see Fig. 6-60).

Several running periods have been devoted to studies with each of the EM calorimeter prototypes. Some of the studies performed are listed below along with a summary of the available results.

- Study of the variation of the response of the calorimeter when the showering particle enters near or on the azimuthal crack between two modules. Data were recorded when the electron had a trajectory parallel to the crack and at angles of 2° , 3° , and 4° with respect to the crack. Data were also recorded for cases in which the crack between the modules was “filled in” with a 1 mm lead sheet. Tentatively, this study shows that there is indeed a dip in the response across module boundaries but that orienting the crack 3° away from the intersection region restores nearly full response. The response as a function of the distance from the azimuthal crack is shown in Fig. 6-61 for the cases of angles of incidence of the electron of 0° and 3° . Insertion of a lead sheet between the two modules to fill the crack decreased the dip when the particle was parallel to the crack, but caused a large




FIG. 6-60. Resolution as a percentage of the mean signal for the EM prototype tile/fiber barrel calorimeter as measured in the Fermilab MP test beam.




FIG. 6-61. Electron shower response shown as a function of vertical position of the incident electron test beam. The response is the summed pulse height of the two towers above and below the azimuthal crack between two EM calorimeter modules. Pulse heights have been normalized to give equal response for each tower for showers in the center of either tower. Beam momentum was 35 GeV and the data have been corrected for the spread in beam momentum.

variation in response. Based on these results we have established a baseline with the crack offset by 3° .

- Study of the linearity of the response from 15–150 GeV for incident electrons. The calorimeter was determined to give a linear response to within 0.5% over this energy range. The data are displayed in Fig. 6-62.




FIG. 6-62. Mean calorimeter response in ADC counts versus electron beam energy. The straight line is a fit to the data for which the measurement at 150 GeV is within 0.5% of the fit value.

- Study of the variation of response of the calorimeter over the area of a single tower. The tower was mapped in a grid of cell size approximately $0.5 \times 0.5 \text{ cm}^2$.
- Measurement of the response of the calorimeter at tower center for varying angles of incidence of the showering particle. This study was performed for angles of 0° and $\pm 3^\circ$.
- Mapping of the calorimeter response in the vicinity of the stainless steel support bulkheads.

The following questions were studied with the scintillating fiber preshower detector:

- Choice of readout. Data were obtained using a CCD coupled to an image intensifier, a multianode photomultiplier tube, and an array of avalanche photodiodes.
- Electron/charged pion separation. Beam line Cerenkov counters were used to tag the beam particles.
- The correction to the measured shower energy with varying numbers of radiation lengths of material before the calorimeter. Data were obtained with a simulated coil of 1 radiation length and with lead inserts of 2 to 6 radiation length thickness.
- Crosstalk in the fiber readout as a function of the voltage on the multianode phototube. Optimization is needed because the relative crosstalk increases with the gain at higher voltages.

The shower maximum detector was embedded in the EM calorimeter prototype and was included in the studies listed above. Different readout methods were tested for the shower maximum detector. Crosstalk and position resolution were studied, as discussed above in Section 6.7.

The hadron prototype partial wedge was installed late in the test beam program and was important in the studies of the preshower detector e/π separation. The hadron calorimeter was studied on its own to study the response across internal boundaries for varying angles of incidence of the showering pion.

References:

1. Letter of Intent by the SDC, Solenoidal Detector Collaboration Note SDC-90-00151, (Nov. 1990).
2. R. N. Cahn *et al.* in Proc. 1987 Workshop on Detectors and Experimental Areas at the SSC, ed. M. Gilchriese (World Scientific).
3. C. Hearty, "Effect of Electromagnetic Calorimeter Thickness on Resolution," Solenoidal Detector Collaboration Note SDC-91-00084, (Sept. 1991).
4. D. Green *et al.*, "Depth requirements of SSC calorimeters," Solenoidal Detector Collaboration Note SDC-91-16 (1991).
5. K. Einsweiler, "Associated Higgs Production in the Intermediate Mass Region," Solenoidal Detector Collaboration Note SDC-91-00095, (Oct. 1991).
6. C. Hearty, "Electromagnetic resolution as a function of lead and scintillator thickness," Solenoidal Detector Collaboration Note SDC-92-185, (Jan. 1992).
7. J. Marrifino *et al.*, "Massless gaps for solenoid plus calorimeter," Fermilab note TN-1766 (1991).
8. T. Kirk, A. B. Wicklund, "Central calorimeter configuration," Solenoidal Detector Collaboration Note SDC-91-18 (1991).
9. C. Hearty, "Impact of the tracking material on EM calorimeter calibration," Solenoidal Detector Collaboration Note SDC-92-216, (Sept. 1991).
10. R. W. Kadel, "Radiation Dose Profiles and Activation of the SDC Endcap Calorimeter," Solenoidal Detector Collaboration Note SDC-92-176, (Jan. 1992).
11. D. Green, A. Para and J. Hauptman, Fermilab FN565, (May 1991).
12. T. A. Gabriel *et al.*, "CALOR89: A Monte Carlo Program Package for the Design and Analysis of Calorimeter Systems," ORNL ORNL/TM-11185.
13. D.E. Groom, "Four component approximation to calorimeter resolution," Capri/91, Solenoidal Detector Collaboration Note SDC-91-37 (1991).
14. T. Kirk, "Shower Counter Resolution Scaling," Solenoidal Detector Collaboration Note SDC-91-108 (1991).
15. P. de Barbaro, *et al.*, "R&D Results on Scintillating Tile/Fiber Calorimetry for the CDF and SDC Detectors," 5th Pisa Meeting on Advanced Detectors, Elba, Italy, (May 30, 1991).
16. B. Lu *et al.*, Preprint VPI-IHEP-91/1, (June 1, 1991).
17. D. Underwood, D. Morgan, J. Proudfoot, "Fiber-Tile Optical Studies at Argonne," Solenoidal Detector Collaboration Note SDC-91-00052 (1991).
18. A. Beretvas and J. Freeman, "Monte Carlo Investigations of Light Yield Uniformity," Fermilab-SDC-100, (Feb. 1992).
19. Q. Fan and S. Olsen, "Effects of longitudinal and transverse variations on the energy resolution of an electromagnetic sampling calorimeter," UR Preprint (1991) (unpublished).
20. D. Kim, "Optical Studies for the CDF Tile/Fiber Calorimeter," CDF-1683.
21. G. Apollinari and S. White, "Plastic Optical Fiber Splicing by Thermal Fusion," Nucl. Instr. and Meth. A311, 520 (1992).
22. R. Wigmans, private communication.
23. V. Barnes *et al.*, "Calibration of scintillating calorimeters," *Proc. of the International Conference on Calorimetry*, (Oct. 29, 1990).
24. C. Zorn *et al.*, Nucl. Inst. Meth. A276, 58 (1989).
25. K. Hara *et al.*, "Radiation Hardness of Scintillating Tile/fibers," Solenoidal Detector Collaboration Note SDC-92-00186 (1992).

26. P. Bonamy *et al.*, "Radiation Damage in Tile-Fiber Calorimeter Modules," Solenoidal Detector Collaboration Note SDC-91-126 (1991).
27. L. Hu *et al.*, "Radiation damage of tile/fiber scintillator modules for the SDC calorimeter," Solenoidal Detector Collaboration Note SDC-91-00119 (1991).
28. P. Bonamy *et al.*, "Radiation Damage in Scintillating Plates and Fibers," Solenoidal Detector Collaboration Note SDC-91-000011 (1991), DPHPE Saclay, Mar 1991.
29. S. Funaki *et al.*, "Beam Test on Radiation Hardness of Scintillating Tile/Fiber Calorimeter," Solenoidal Detector Collaboration Note SDC-91-00085 (1991).
30. Y. Kikuchi, "Performance of the Electromagnetic Shower Counter for the Collider Detector at Fermilab," Master's Thesis (1985) University of Tsukuba (unpublished).
31. ZEUS collaboration, "Detector Status Report," 1989.
32. *Proc. of Workshop on Radiation Hardness of Plastic Scintillator*, FSU, Mar 1990.
33. V. Hagopian *et al.*, "Radiation Damage Tests of New Scintillating Fibers and plates," Symposium on Detector Research and Development for the Superconducting SuperCollider (1990).
34. A. R. Fazely *et al.*, "Radiation Damage Studies for the SDC Electromagnetic Calorimeter," Solenoidal Detector Collaboration Note SDC-92-172 (1992).
35. V. Guarino *et al.*, "Mechanical Design and Finite Element Analysis of the SDC Central Calorimeter," Proc. IISCC New Orleans 1992, ANL-CP-92-13.
36. P. Bonamy, *et al.*, "Preshower and Shower-Maximum Detectors for SDC," Solenoidal Detector Collaboration Note SDC-90-00153, (Dec. 1990).
37. D. Muller, *Phys. Rev.* **5**, 2677 (1972).
38. P. Bonamy, *et al.*, "Radiation Damage in Tile-Fiber Calorimeter Modules," Solenoidal Detector Collaboration Note SDC-91-126, (Nov. 1991).
39. P. Bonamy, *et al.*, "Radiation Damage in Scintillating Plates and Fibers," Solenoidal Detector Collaboration Note SDC-91-000011, (Mar. 1991).
40. V. Barnes, private communication.
41. F. Abe *et al.* (CDF), *Phys. Rev.* **D43**, 2070 (1991).
42. F. Abe *et al.* (CDF), "The Topology of Three Jet Events in $p\bar{p}$ Collisions at $\sqrt{s} = 1.8$ TeV," *Phys. Rev.* **D45**, 1448 (1992).
43. S. Hahn *et al.*, *Nucl. Instr. and Meth.* **A267**, 351 (1988).
44. N. Hill, V. Guarino, J. Nasiatka, M. Burke, R. Swensrud, *proc. IISCC Atlanta 1991*, M. A. Burke *et al.*, *proc. IISCC New Orleans 1992*.
45. Contributions by Rick Field and by Mike Barnett to the Forward Calorimeter Workshop, held at LBL, 10/31-11/01, (1991).
46. L. Demortier *et al.*, "Radiation Damage Studies on High Pressure Gases," Solenoidal Detector Collaboration Note SDC-92-182 (1992).
47. L. Demortier *et al.*, "Test Beam Results of a Prototype Parallel Plate Electromagnetic High Pressure Gas Calorimeter," Solenoidal Detector Collaboration Note SDC-92-183 (1992).
48. L. Demortier *et al.*, "A High Pressure Gas Ionization Tube Calorimeter for the Forward Region," Solenoidal Detector Collaboration Note SDC-92-184 (1992).
49. M. Chiba *et al.*, Conference on solenoidal detectors at colliders, Tsukuba (1990).

7. Muon system

7.1. Introduction and overview

The muon system has three distinct functions:

1. To trigger the detector on a muon above a threshold p_t ,
2. To identify a charged track as a muon, and
3. To provide a muon momentum measurement in conjunction with the inner tracker.

Each of these functions makes different demands on the muon system, and, in general, each part of the muon system participates in more than one function. In what follows, we briefly describe the components of the muon system, and then explain how the various parts work to satisfy the system requirements. Later sections of this chapter describe the various components in greater detail.

7.1.1. The components

A schematic of the muon system has been shown in Fig. 2-7. It divides logically into the central region, which includes the barrel and intermediate systems, and the forward region. (For simplicity of notation, references to a barrel system will implicitly include the corresponding intermediate system.) The muon system has three main components: magnetized iron toroids, wire chambers, and trigger counters.

The central toroid is 1.5 m thick and covers the polar angle region $28^\circ < \theta < 152^\circ$, or $|\eta| < 1.4$. When magnetized to 1.8 T, it provides a polar angle (θ) kick of $810/p_t$ mr, where p_t is measured in GeV/c. There are two forward toroids at each end of the detector, and each toroid is 1.5 m thick. They cover the angular regions of $9.4^\circ < \theta < 28^\circ$ and $152^\circ < \theta < 170.6^\circ$, corresponding to $1.4 < |\eta| < 2.5$. The polar angle kick from the forward toroids is $1620 \tan \theta / p_t$ mr. This corresponds to $860/p_t$ mr at $\theta = 28^\circ$ and to $268/p_t$ mr at $\theta = 9.4^\circ$.

The wire chambers are of a novel design. In the central region ($|\eta| < 1.4$), they consist of 9.32 cm outer diameter round aluminum tubes glued together on thin aluminum plates to make a rigid structure. The ends of the anode wires are precisely positioned by being referenced to a CNC-milled plate. Each tube contains a pair of electrodes to shape the electric field so that electrons drift in a nearly uniform field. Each tube is rotated so that the field lines are perpendicular to the trajectory of an infinite momentum particle originating at the interaction point. This field orientation allows the chambers to have good position resolution and two-track separation. The chambers in the forward region are of a similar design, but are 4.5 cm outer diameter in layers before the forward toroids, FW1 and FW2, and 6.0 cm outer diameter in layers after the toroids, FW4 and FW5.

The basic superlayer of chambers consists of four layers of tubes. Wires in the first and third layers, and in the second and fourth layers, are projective to the interaction region, and the two sets of layers are offset by approximately one half cell from each other. The projective wire arrangement is used to form a simple trigger, as will be explained below.

In the central region, wire chambers are positioned to measure in both the polar angle (θ) and azimuthal angle (ϕ) directions before and after the toroid. In addition, two single layers of tubes with small angle stereo to the θ direction are located in BW3. In the forward region, chambers are positioned to measure in the θ direction and in small angle stereo ($\pm 7.5^\circ$) to the θ direction.

Scintillation counters are used in the Level 1 trigger to define the beam crossing time of the interaction. The central region contains a single layer of counters, which are placed immediately after BW2. Each counter is approximately 0.5 m wide and 2.0 m long, with the long axis parallel to the θ -measuring chambers. It is read out by two photomultipliers, each viewing one end to provide a mean-time signal.

In the forward regions, two layers of scintillation counters are placed after FW4 and FW5. In this region each counter is read out by a single photomultiplier.

Čerenkov counters, which are insensitive to neutrons and low-momentum charged particles, can be inserted between FW4 and FW5 to assist in the Level 1 trigger if needed. This option will be discussed below, but is not part of the baseline design.

The detailed layouts of the various components in the barrel and forward regions are given in Tables 7-1 and 7-2.

Table 7-1
Layout of muon system components in the barrel region.

Label	Type	Coordinate	Layers	Inner r (m)
BW1	Wire chamber	θ	4	5.76
		ϕ	4	6.16
BT1	Iron toroid			6.75
BW2	Wire chamber	θ	4	8.48
BS2	Scintillator	θ	1	8.91
BW3	Wire chamber	θ	4	9.96
		ϕ	4	10.36
		s	2	10.76

Table 7-2
Layout of muon system components in the forward regions.

Label	Type	Coordinate	Layers	Inner z (m)
FW1	Wire chamber	θ	4	7.54
FW2	Wire chamber	θ, s_1	4	9.17
		θ, s_2	4	9.75
FT1	Iron toroid			10.16
FW3	Wire chamber*	θ	4	11.75
FT2	Iron toroid			12.27
FW4	Wire chamber	θ	4	13.86
FS4	Scintillator	θ	1	15.52
FC4	Čerenkov*		1	15.72
FW5	Wire chamber	θ, s_1	4	18.82
		θ, s_2	4	19.27
FS5	Scintillator	θ	1	19.72

* Possible future upgrade. Not part of the baseline design.

7.1.2. The trigger

The Level 1 trigger operates at the beam crossing rate of 60 MHz and must arrive at a decision and transmit it to all components of the detector within the 4 μ s pipeline length. Thus, it must be simple and fast, and still have the ability to reduce the accepted rate for each detector component to a few kHz. The Level 1 muon trigger is formed as a coincidence between scintillation counters and wire chambers which lie along a road in the θ direction.

The key to this trigger is a local measurement of the bend of a muon candidate in the toroids. This is done by measuring the difference in time between signals from projective wires. An infinite momentum muon would have zero time difference if it originated at the center of the interaction region. Finite momentum tracks have non-zero time differences because they bend in the toroid, multiple Coulomb scatter in the calorimeter and the toroid, originate away from the center of the interaction region, and are measured with finite resolution by the chambers. Table 7-3 contains values for these contributions at three values of η .

Table 7-3

Contributions to the time difference measured by the Level 1 trigger at three values of η . Values are given in terms of the angle difference in mr. p_t is measured in GeV/c. The value for the interaction point uncertainty is for the innermost relevant measuring plane.

η	θ (degrees)	Bend	Scattering	Interaction point	Chamber resolution
0.0	90	$810/p_t$	$210/p_t$	5.9	1.8
1.0	40	$810/p_t$	$160/p_t$	2.5	1.2
2.0	15	$450/p_t$	$70/p_t$	0.9	1.7

We would like the Level 1 trigger to be efficient for single muons with p_t above 20 GeV/c, and for multi-leptons with p_t above 10 GeV/c. The major contribution to the trigger rate is lower p_t particles, primarily real muons, which manage to satisfy the trigger. One way this can happen is by a particle crossing on different sides of a projective wire pair. This is illustrated in Fig. 7-1. Simulations have shown that this problem can be eliminated by requiring signals from two or more pairs of projective wires along a road in θ to satisfy the trigger. Since there are four sets of projective wires, two each in BW2 and BW3 (FW4 and FW5 in the forward region), this requirement yields a trigger with high efficiency. Figure 7-2 shows typical Level 1 trigger efficiency curves for various time difference thresholds.

The requirement that a muon pass through a projective pair of drift tubes places a minimum size requirement on the drift tubes. This is mainly relevant in the forward region where drift tubes of at least 6 cm diameter must be used in FW4 and FW5 to obtain full efficiency for a p_t threshold of about 10 GeV/c.

It is clear from Table 7-3, that, with a steeply falling muon momentum spectrum, the major contributor to the Level 1 muon trigger rate will be below-threshold muons that scatter in the calorimeter and toroid in such a way as to appear to have higher momenta. Estimates of the rate for the Level 1 muon trigger are of order 6 kHz, which is near the limit of acceptability. However, these are difficult calculations, which have considerable uncertainty [1]. It may well be necessary to augment the basic Level 1 muon trigger with additional components. Possibilities include the following:

1. The requirement that there be a high- p_t θ stub in BW1 (FW1 or FW2) pointing to the triggering chambers in BW2 and BW3 (FW4 and FW5). This requirement would eliminate low- p_t muons which had large angle scatters in the calorimeter.
2. The requirement that there be a stub in BW1 (FW1 or FW2) pointing to the triggering chambers in BW2 and BW3 (FW4 and FW5) such that the angle measured before and after the toroid is consistent




FIG. 7-1. An illustration of a low- p_t muon satisfying the Level 1 trigger on a single pair of projective wires by passing on different sides of the wire. Note that this track does not satisfy the trigger in the set of cells offset by a half cell, shown as dashed circles.




FIG. 7-2. Trigger efficiency versus transverse momentum for three different time difference thresholds.

with expectations. (The difference between this requirement and the previous one is that here one must correlate two angle measurements, whereas in the previous requirement it is only necessary to find a stub in BW1 pointing to the interaction point in a road which matches the BW2-BW3 trigger. Clearly the former requirement is much easier to implement than the latter.)

3. The requirement that there be a high- p_t stub in BW1 or BW3 in the ϕ region indicated by the triggering scintillation counter. (By "high- p_t stub," we simply mean one that points to the interaction region within the multiple scattering error.) This requirement is useful for reducing the expected Level 1 trigger rate from cosmic rays to a negligible level. Without the ϕ trigger, the additional Level 1 trigger rate is estimated to be about 1 kHz [2].
4. The requirement that there be a track in the inner tracker of sufficient p_t in the ϕ region indicated by the triggering scintillation counter.
5. The requirement that the calorimeter towers pointed to by θ and ϕ measurements in the muon system be quiet. This requirement may allow a lower threshold for isolated muons, but is not useful for a general muon trigger.

The function of the Level 2 trigger is to refine the Level 1 trigger by sharpening the p_t threshold. In the central region, the primary method is to match a track in the inner tracker to a ϕ measurement in BW1 or BW3. An angle-angle θ measurement in the muon system, such as that discussed above as an option for the Level 1 trigger, may also be of help. In the forward region, where the thicker toroids provide more resolution, this is the primary method. It is accomplished by a measurement of the bend in the toroid using FW1 and FW2 to establish the direction before the toroid and FW4 and FW5 to establish the direction

7.3.1. Central region

Dimensions and channel count

The exterior area of the barrel toroid is covered with a single layer of rectangular scintillation counters whose width, along the z direction, is 0.50 m and whose length is approximately 1.85 m. A total of 2240 counters cover this area, 1920 in BS2 and 320 in IS2. The length of each scintillation counter is such that four counters span the width of an octant, providing an azimuthal segmentation which matches that of the calorimeter. Two fundamental factors prevent the use of longer counters. First, typical light attenuation lengths in scintillator are around 2 m, setting this as a reasonable scale for the length of counters. Second, and even more important, the variation of light transit times, corresponding to an effective light speed in scintillator of 16 cm/ns, must be kept well below the beam crossing period of 16 ns to assure the correct beam tagging. As described below, mean-timing methods will be used to give a precise measure of the scintillator signal arrival time. However, the system must function even when the muon is accompanied by showers of electromagnetically produced particles that may confuse the mean-timer.

The scintillator thickness and width are set by light yield and collection criteria. Based on measurements shown in Fig. 7-16, a thickness of about 10 mm appears to provide a sufficient photoelectron count for most choices of scintillator. Our goal is a minimum of 20 photoelectrons from any region of a counter to ensure efficient muon detection even if there is some degradation of scintillator response with time.




FIG. 7-16. Measurements of photoelectron yields in various scintillators as a function of distance from the light guide. The far end of each scintillation counter was covered with black paper.




FIG. 7-17. Measured photoelectron yields from a 0.60 m wide counter as a function of distance from a twisted-strip light guide for direct transmission and for light reflected by a 180° reflector.




FIG. 7-3. Muon momentum resolution from measurements in the θ direction, assuming that the event vertex is accurately known in x , y , and z . The resolution is shown for $p_t = 0.1, 0.3, 1.0$, and 3.0 TeV/ c from the lowest to highest curve.

discard low- p_t muons that fake high- p_t ones, as well as to recover from an incorrect momentum measurement of high-energy muons.

7.1.4. Momentum measurement

In the SDC detector, the primary momentum measurement for muons comes from the inner tracker and the solenoidal field. However, for $p_t > 300$ GeV/ c , the highest precision measurements of muon momentum in the central region come for a combination of ϕ measurements in the inner tracker and in BW1 and BW3 assuming we can monitor accurately the relative alignment between the central tracks and muon system. This is because the effect of the solenoidal bend extends out to the point in the calorimeter at which all of the magnetic flux has been returned. At very forward angles, the toroidal measurement improves the precision of the solenoidal measurement for $p_t > 100$ GeV/ c and becomes more precise than the solenoidal measurement around a p_t of 300 GeV/ c . Figure 7-4 shows the muon momentum resolution as a function of η for the inner tracker alone, and for the combined system, where again accurate knowledge of the event vertex position has been used.

The main degradation in the absence of precise vertex information occurs at forward rapidity. A muon with $p_t = 1$ TeV/ c has a resolution at $\eta = 2.5$ of 18% with vertex information and 30% in the total absence of inner tracker and vertex information. This difference relates to whether the muon direction upstream of the forward toroids is defined by a line drawn from the event vertex to FW2, or by a line drawn from FW1 to FW2.

At higher-than-design luminosities, the forward muon system may be required to perform all of its functions independently of the inner tracking system. The present design allows this, albeit with a reduced momentum resolution, as shown in Fig. 7-5.

7.2. Design choices

In designing the muon system, a number of choices had to be made. The driving force behind these decisions was to maximize physics capabilities while minimizing cost. The capabilities we considered were the momentum resolution of the system, including the highest momentum at which a sign determination can be made, trigger resolution and rates, track matching to the central tracker, and the robustness of the system. In this section, we review some of the considerations that went into the design choices.




FIG. 7-4. Momentum resolution for muons using combined measurements from the inner tracker and the muon system (solid lines) and from the inner tracker alone (dashed lines).




FIG. 7-5. Momentum resolution in the forward region based solely on the muon system.

7.2.1. Toroid thicknesses

The central toroid thickness was chosen to be 1.5 m. We considered a 1.0 m-thick toroid for some time and rejected it as too risky. One consideration had to do with the resolution of the momentum measurement. In the measurement dominated regime (high p_t), the resolution clearly becomes 50% worse for a 1.0 m-thick toroid. In the multiple scattering dominated regime, it becomes 22% worse, since the multiple scattering varies as the square root of the thickness.

However, a more serious concern was the effect on the basic Level 1 trigger. As can be seen from Table 7-3, at $\eta = 0$ the ratio between the toroidal bend and the rms multiple scattering is 3.9 for a 1.5 m-thick toroid. This ratio changes to 2.8 for a 1.0 m toroid, since the multiple scattering is dominated by the calorimeter in this case. Since the Level 1 trigger rate is dominated by low-momentum muons which scatter in such a way as to appear to have a higher momentum, this would provide a substantial increase in the Level 1 trigger rate. We have noted above that this rate is already somewhat large.

One way to control the Level 1 trigger rate is to increase the p_t threshold of the trigger. Our ability to do this is reduced with a 1.0-m thick toroid because the finite size of the beam crossing diamond, nominally 5 cm rms, limits the p_t at which a Level 1 trigger can be set. For sake of illustration, assume that the toroidal bend must be at least twice the effect of the rms longitudinal beam length for the p_t threshold to be sensible — above this threshold the trigger rate would not decrease appreciably for an increase in the p_t threshold. From Table 7-3, this sets the maximum p_t threshold at 69 GeV/c for a 1.5 m-thick toroid and at 46 GeV/c for a 1.0 m-thick toroid. If the SSC should be run with twice as long a beam diamond, then these values become 34 and 23 GeV/c, respectively. The latter case provides almost no leeway in setting the p_t threshold of the trigger. Figure 7-6 illustrates the above argument by showing a typical trigger efficiency versus p_t for two toroid thicknesses.

The forward toroid thickness is determined by the goal of having good momentum resolution and triggering capability in the high rapidity region in which the inner tracker measurement is weaker. Examination of Fig. 7-4 shows that for $|\eta| > 2$ and $p_t > 100$ GeV/c, the forward muon system is significantly improving the momentum resolution. In fact, for $p_t \geq 300$ GeV/c, the resolution at $|\eta| = 2.5$ is less than a factor of two worse than it is at $\eta = 0$. We believe that good lepton energy resolution over the rapidity interval up to $|\eta| = 2.5$ is an essential requirement for a general-purpose detector at the SSC.

The muon trigger is more difficult in the forward direction because of the lack of direct ϕ readout. The resolution from the toroidal measurement must be sufficient to provide an acceptable Level 2 trigger rate.

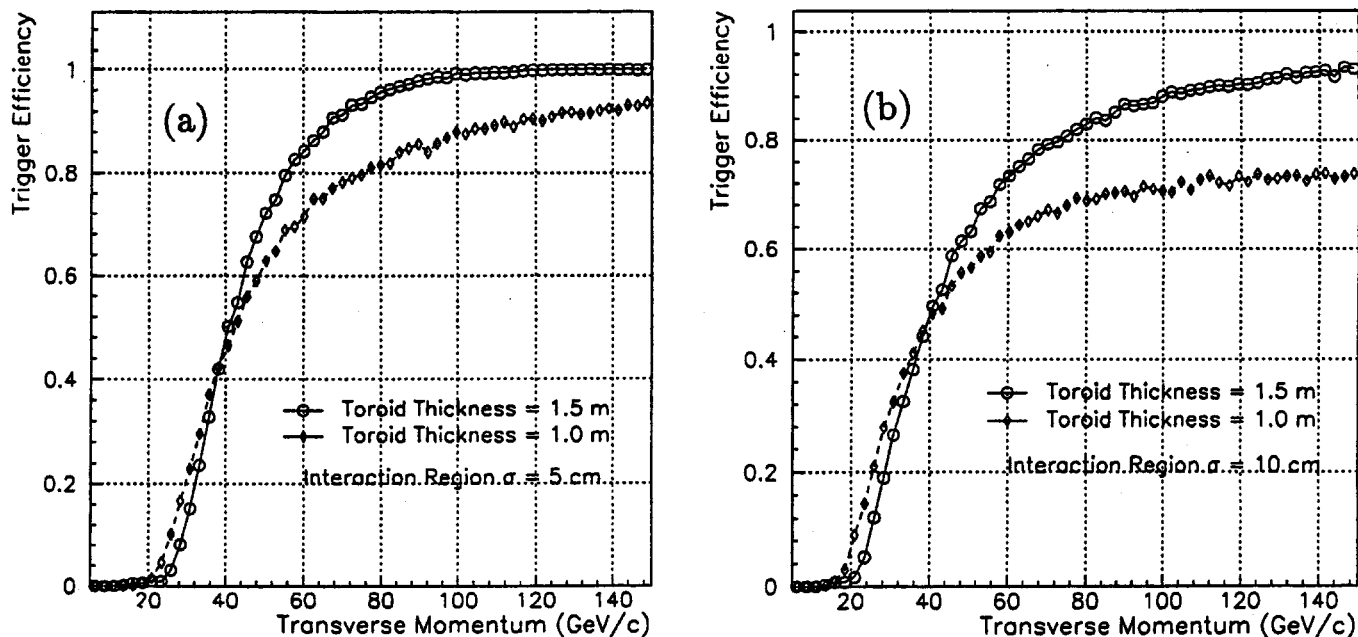


FIG. 7-6. Trigger efficiencies versus p_t , for a nominal 40 GeV/c threshold, for 1.5 and 1.0 m-thick central toroids, assuming (a) the nominal $\sigma = 5$ cm beam interaction diamond and (b) a $\sigma = 10$ cm diamond.

Iron thicknesses of less than 3 m will not allow this.

7.2.2. Chamber resolution

The natural resolution of drift tubes with shaped fields of the size being used here is about $250 \mu\text{m}$, whereas drift tubes without shaped fields tend to have resolutions of about $500 \mu\text{m}$ [4]. For a $250 \mu\text{m}$ measurement resolution to be used effectively, it must be coupled with an alignment system which can track and determine positions to about $150 \mu\text{m}$. Figures 7-7 and 7-8 show the effect of doubling the chamber resolution to twice the design value ($500 \mu\text{m}$ per wire and $300 \mu\text{m}$ alignment) for the combined system and for the muon system momentum measurement, respectively. The combined system resolution is sensitive to this change only for $p_t > 1$ TeV/c, while the momentum measurement resolution from the muon system only is affected for p_t 's ≥ 100 GeV/c; it is about a 30% effect at $p_t = 1$ TeV/c.

There is little, if any, cost advantage in using drift tubes without field-shaping electrodes. Drift tubes without field shaping also have considerably worse track-pair separation. The reason for having good two-track separation is discussed next.

7.2.3. Number of layers

The number of layers has been carefully determined to be the minimum needed for a robust system. In this regard, it is important to remember that very high-energy muons are often accompanied by electromagnetic showers. Recent test beam results, as well as GEANT simulations, indicate that when several hundred GeV muons emerge from material, they are accompanied by extra particles 20 to 30% of the time [5]. Figure 7-9 shows a typical event of this type from the recent Fermilab T-816 test. The fraction of muons with accompanying debris is highest when the muons have traversed a large amount of iron, but even a few centimeters of aluminum, typical of chamber walls, is sufficient to generate a fair amount of debris.

The debris tends to be at large angles to the muon, and normally is not a problem for chambers with good two-track separation, such as the ones we are proposing. However, some of the time, a measurement will be corrupted by the wrong particle creating the signal. For this reason, we have designed the system so that it does not rely on a single superlayer for a critical measurement.

The reason for using four layers to form a basic superlayer is clear. The first and Level 2 triggers use projective wires in both coordinates to select stiff tracks. One pair of projective wires is not sufficient since




FIG. 7-7. Muon resolution with the nominal muon chamber position resolution (dashed lines) and with muon-chamber resolution a factor of two worse (solid lines). Measurements from the inner tracker and muon system are combined.




FIG. 7-8. Same as Fig. 7-7 but for momenta measured in the muon system alone, assuming that the event vertex is accurately known.

run22.dat

SR#

15 EV#

16




FIG. 7-9. A typical event from Fermilab T-816 in which a high-energy muon is accompanied by other particles. Due to left-right ambiguities each signal is shown twice. Only the vertical measuring chambers are shown. The total vertical height is 21 cm. The figure is not to scale, with longitudinal gaps of 1 m between each of the chambers not shown. There is 64 cm of iron between the second and third chambers. There are two incident muons; the second one is difficult to see since it is exactly at the boundary between two cells. The wide-angle nature of the debris is seen by the observation that the final chamber has only four tracks, the initial two muons plus two other particles, while the third chamber has many tracks.

muons passing near the walls of the tube chambers can easily miss one of the pair of chambers. However, the approximate half-cell offset between the pairs of projective wires in a superlayer guarantees that at least one pair is efficient for the passage of a moderate p_t muon.

There are two superlayers of θ chambers after the central toroid and one superlayer before. The outer superlayers form the basic Level 1 trigger and are thus needed for robustness. The reason why at least two sets of projective pairs are necessary for the trigger was explained in Section 7.1.2. Having four pairs

available allows a two out of four majority to be formed.

There will be cracks in any superlayer where modules come together. Cracks in BW2 and BW3 along the beam line are not projective, so that either superlayer can complete the trigger. In the baseline design, the cracks in the azimuthal direction are projective, but a modification to eliminate projective cracks is under study.

The BW1 θ superlayer is used in the redundant momentum measurement, may be used to augment the Level 1 trigger, and is useful for track-matching to the inner tracker. It has no redundancy, but its functions are not as critical as those of the other superlayers. For example, if a measurement is missing in BW1, the redundant momentum measurement can be done using the inner tracker to determine the incoming direction, at the cost of increased error due to multiple scattering in the calorimeter.

Azimuthal (ϕ) measurements are critical for the Level 2 trigger, possible augmentation of the Level 1 trigger, track-matching to the inner tracker, and high-precision momentum measurements for very high p_t muons. Two superlayers are clearly needed. We have placed one inside the central toroid, where there is less multiple scattering, but a more hostile environment from calorimeter leakage, and one outside the toroid, where the environment is favorable, but there is more multiple scattering due to the toroid. The effect of ϕ chamber placement on track matching is discussed in the next subsection.

The central region is completed by two single layers of small-angle stereo chambers, a bare minimum.

Most of the above arguments apply to the forward system also. The forward system differs from the central in two main respects. First, it has a better and more independent momentum measurement, motivated by the fact that the solenoidal measurement is weaker. Second, it uses small angle stereo instead of ϕ measurements. There are two reasons for not using ϕ measurements in the forward direction: it is technically difficult, and it is unwise because it links measurements at very different values of η . High occupancy at high η can affect measurements at lower η .

The two superlayers FW1 and FW2 before the forward toroids ensure a good θ measurement upstream of the toroid even in the absence of precise inner tracker information.

7.2.4. ϕ chamber placement

As noted above, one of the functions of the ϕ chambers is to match signals in the muon system to tracks found by the inner tracker. We define $\Delta\phi$ as the azimuthal region that must be searched to find a match for a projected track from the central tracker. Figure 7-10 shows the dependence of $p_t \times \Delta\phi$, which we call the "confusion volume," on the muon p_t for the cases of placing ϕ chambers in BW1, BW2, and BW3, or all three. Because $\Delta\phi$ is mostly determined by multiple scattering, the product $p_t \times \Delta\phi$ has a much weaker p_t dependence than $\Delta\phi$ itself. There is a clear advantage to having ϕ measurements in BW1.

The four ϕ layers in BW1 provide measurements of both the ϕ angle and the intercept of muons multiply scattered from the barrel calorimeter. The strong correlation between the scattering angle and the independently measured intercept can be exploited to reduce the loss of information due to multiple scattering and thus significantly improve the ability to match tracks [3].

7.2.5. BW2-BW3 and FW4-FW5 separation

Figure 7-11 shows the resolution of the muon system momentum measurement in the barrel region as a function of the radial gap between BW2 and BW3, which is about 1 m in the baseline design. There is a loss of resolution if the gap is removed, but there is little advantage in making it larger. We note that the gap is a convenient place to locate the scintillation counters. It should thus be large enough for personnel access to service the scintillation counter photomultipliers.

Figures 7-12 and 7-13 show the resolution of the combined system momentum measurement and the muon system momentum measurement in the forward region as a function of the distance between FW4 and FW5, 4.3 m in the baseline design. The knees of the curves are near the proposed value, which implies there is a good match to the chamber resolution.




FIG. 7-10. The "confusion volume," defined in the text for matching a track in the muon system to a track in the inner tracker for a system which has ϕ -measuring chambers in BW1 only (dashed line), BW3 only (dash-dotted line), and both, as in the present design (solid line). Both position and angle measurements in the system are important, but their effectiveness varies with p_t , causing the shoulders in the curves.




FIG. 7-11. The effect of changing the gap between BW2 and BW3 on the muon system momentum measurement resolution at $\eta = 0$. The abscissa plots the difference between a given gap and the baseline value (92 cm). The resolution is shown for $p_t = 0.1, 0.3, 1.0$, and 3.0 TeV/c from the lowest to the highest curve.




FIG. 7-12. The effect of changing the z distance between FW4 and FW5 on the combined inner tracking and muon system momentum resolution at $|\eta| = 2.4$. The ordinate plots the difference between a given distance and the proposed value (4.3 m). The resolution is shown for $p_t = 0.1, 0.3, 1.0$, and 3.0 TeV/c from the lowest to the highest curve.




FIG. 7-13. The effect of changing the z distance between FW4 and FW5 on the muon system momentum measurement resolution alone at $|\eta| = 2.4$, assuming the event vertex is accurately known. The ordinate plots the difference between a given distance and the proposed value (4.3 m).

7.3. Scintillation counters

The SDC muon system employs scintillation counters as the primary means of identifying the beam crossing time associated with each detected muon. The wire chambers are unsuitable for this function since their maximum drift time of about $1.0 \mu\text{s}$ corresponds to 60 beam crossings. Identification of the correct beam crossing is essential at all trigger levels and also in the offline analysis since event reconstruction combines data from many detector elements that must be properly correlated in time. The required time resolution must allow an unambiguous tagging of the beam crossing, since an average of 1.6 inelastic events are generated per crossing at the SSC design luminosity. The resolution requirements are set by the beam crossing period of 16 ns.

The scintillation counters do not by themselves provide a complete muon trigger but are always operated in coincidence with signals generated by the muon drift chambers. Trigger roads are defined by appropriate clusters of drift cells in BW2 and BW3 (or IW2-3, FW4-5) and their overlapping scintillation counters.

One layer of scintillation counters is sufficient in the central region of the detector, outside the toroid, to provide a good time measurement at the SSC design luminosity. Each counter in the central region will be instrumented with two photomultipliers to provide a redundant readout. At smaller angles, in the region behind the forward toroids, two separated layers of scintillation counters will be used in the trigger to suppress localized backgrounds from low energy particles generated in various masks and shields around the beam pipe and focusing quadrupoles. Each counter in the two-layer system will be read out by a single photomultiplier, and appropriately matched counters from the two different layers will be placed in coincidence. The general configuration of counters in the barrel and forward regions is shown in Figs. 7-14 and 7-15, respectively.




FIG. 7-14. Configuration of muon scintillation counters in the barrel region.




FIG. 7-15. Configuration of muon scintillation counters in the forward region.

Backgrounds at small angles are especially difficult to predict since they depend strongly on the exact details of the detector and beamline configuration. For this reason, we propose as an upgrade option the addition of Čerenkov counters in the small-angle region just behind the forward muon toroids. Čerenkov counters are relatively insensitive to neutron backgrounds and have a highly directional response. (Čerenkov counters will be discussed in the next section.) The muon trigger system has therefore the following upgrade possibilities in case of unforeseen backgrounds or operation at very high luminosities: at small angles, the addition of Čerenkov counters, and at large angles, the addition of a second layer of scintillation counters.

7.3.1. Central region

Dimensions and channel count

The exterior area of the barrel toroid is covered with a single layer of rectangular scintillation counters whose width, along the z direction, is 0.50 m and whose length is approximately 1.85 m. A total of 2240 counters cover this area, 1920 in BS2 and 320 in IS2. The length of each scintillation counter is such that four counters span the width of an octant, providing an azimuthal segmentation which matches that of the calorimeter. Two fundamental factors prevent the use of longer counters. First, typical light attenuation lengths in scintillator are around 2 m, setting this as a reasonable scale for the length of counters. Second, and even more important, the variation of light transit times, corresponding to an effective light speed in scintillator of 16 cm/ns, must be kept well below the beam crossing period of 16 ns to assure the correct beam tagging. As described below, mean-timing methods will be used to give a precise measure of the scintillator signal arrival time. However, the system must function even when the muon is accompanied by showers of electromagnetically produced particles that may confuse the mean-timer.

The scintillator thickness and width are set by light yield and collection criteria. Based on measurements shown in Fig. 7-16, a thickness of about 10 mm appears to provide a sufficient photoelectron count for most choices of scintillator. Our goal is a minimum of 20 photoelectrons from any region of a counter to ensure efficient muon detection even if there is some degradation of scintillator response with time.




FIG. 7-16. Measurements of photoelectron yields in various scintillators as a function of distance from the light guide. The far end of each scintillation counter was covered with black paper.




FIG. 7-17. Measured photoelectron yields from a 0.60 m wide counter as a function of distance from a twisted-strip light guide for direct transmission and for light reflected by a 180° reflector.

Light collection and readout

The barrel counters are read out at both ends using standard photomultipliers with a diameter of 50 mm. The requirement of coincident signals from both ends eliminates spurious trigger signals from tube noise. It also allows the implementation of mean-timing techniques which sharpen the time resolution from the 12 ns variation in light transit times to about 1 ns.

The most efficient light collection scheme is provided by a twisted strip plastic light guide which channels photons adiabatically into the photomultiplier. A guide with four or five 50 mm wide strips is likely to be optimal in terms of cost and performance. One potential problem associated with the use of such light guides is the generation of Čerenkov light by muons or other charged particles. In a close-packed one-layer scintillation counter system, all light guides overlap the scintillator of adjacent counters so that a single muon can, in principle, give multiple hits: from the scintillator in one counter and from the overlapping light guide of the adjacent counter. The fraction of geometric overlap is about 10–20% of the total area, depending on the details of the light guide design. Preliminary measurements and analysis show this to be a negligible problem. The Čerenkov signal generated in the light guide is small, typically of order 5–10 photoelectrons in the photomultiplier next to the light guide. This signal is attenuated by a factor of about three at the other end of the counter and will generally be too small to produce a coincident pulse above threshold. However, even if the Čerenkov signal is detected at both ends, the mean-timer will produce a signal that is synchronous with that from the overlapped scintillator. Hence, no spurious beam tagging will occur. The case in which a single muon is accompanied by a large electromagnetic shower is more complicated. The light guides represent an effective lengthening of the counter and may lead to a slight spill-over of triggers into adjacent, incorrectly identified beam crossings. This case must still be investigated with a detailed simulation program, but its occurrence is estimated to be small (of order 1% or less).

Light collection at the edge of an octant presents a special challenge since scintillation photons must be reflected by 180° at this location. As shown in Fig. 7-17, a triangular prism with aluminized reflecting surfaces transmits about 40% of light exiting the end of the scintillator. This figure shows that we can reasonably expect more than 20 photoelectrons from a 50 cm wide counter even for the worst case of a simple 180° reflector. We expect improved performance with an adiabatic 180° light guide to be tested in the near future. (Preliminary results indicate that approximately 60% of the light can be transmitted through a compact, curved 180° reflector.)

The primary function of the counters is to provide a precise measurement of the traversal time of muons. Since the light is read out at opposite ends of the counters, mean-timing methods allow one to remove the effect of the finite propagation speed of light in the counter. Figure 7-18 shows measurements made with a mean-timer of cosmic rays traversing a 2.7 m long counter at three widely separated locations. The mean-timer output is essentially independent of muon position with a resolution of less than 1 ns after unfolding the time jitter of the cosmic-ray trigger counters providing the start signal.

Rates in the barrel

A number of studies have been made of particle rates outside the barrel toroid [6]. Different sources and their rates are summarized as follows, where the rate is calculated for the interval $-1.5 < \eta < 1.5$ and assumes an SSC luminosity of $10^{33} \text{ cm}^{-2}\text{s}^{-1}$ [7]:

(a) Muons from pion and kaon decay	42 kHz
(b) Muons from charm and bottom decay	45 kHz
(c) Cosmic-ray muons	20 kHz
(d) Punch-through and detected low-energy neutrons	< 100 kHz

Calculations for (a)–(c) give reasonably consistent results and are expected to be accurate to within a factor of two. Punch-through and neutron rates are more difficult to estimate, particularly the latter. A rate of 100 kHz represents a conservative upper limit; the actual rate may be substantially lower, perhaps at the 10 kHz level. We thus estimate an upper limit of 200 kHz spread fairly uniformly over the approximately




FIG. 7-18. Mean-timer output of a 2.7 m long scintillation counter read out at both ends. The variation in output times includes the effect of time jitter of small cosmic-ray trigger counters that provide the start signal.

2000 counters of the central region. This yields a rate of 100 Hz per counter, or 800 Hz for a trigger segment of eight counters, corresponding to two adjacent strips spanning an octant. These low rates indicate that the probability of uncorrelated, extra hits in a trigger segment during the $1.0 \mu\text{s}$ drift time of the muon chambers will be below the 0.1% level at the SSC design luminosity.

7.3.2. Forward region

Dimensions and channel count

The region behind the forward muon toroids is covered with two widely separated layers of scintillation counters, FS4 and FS5. Each counter in these layers is read out by a single photomultiplier. The two layers are segmented and aligned such that a ray, extending from the nominal interaction point through the center of a counter in the first layer, strikes the boundary between two counters in the second layer. A trigger signal is then defined by the coincidence of a signal from a counter in the first layer with a signal from the logical OR of these two counters in the second layer. This particular arrangement allows the definition of appropriate trigger roads in θ , the coordinate in which the magnetic bending of muon trajectories takes place, without resorting to overlapping scintillation counters in the second layer. It is desirable that these trigger roads define a transverse momentum response that is independent of angle. This is achieved by scaling the widths of counters, along the radial coordinate, as $\sin^{1.35} \theta$.

Table 7-4 summarizes the geometry of the first forward layer, FS4, while Fig. 7-19 shows the trigger efficiency as a function of muon transverse momentum. There are 70 counters in the first layer and 71 counters in the second layer of each octant, giving a total of 2256 counters for both ends of the complete forward system. We note that the dimensions of counters in Table 7-4 are provisional and depend on the final configuration of detector elements in the forward region. In particular, the final design will approximate the scaling of counter widths with a smaller set of discrete dimensions to simplify construction. This is not expected to affect the channel count significantly.

Table 7-4
Dimensions of the first layer of forward muon trigger counters.*

Radial boundaries for regions which have 2 ϕ segments per octant (in m)	2.500	2.570	2.643	2.718	2.796	2.877
	2.961	3.048	3.139	3.233	3.331	3.433
	3.539	3.649	3.763	3.882	4.005	4.134
	4.268	4.407	4.552	4.703	4.861	5.024
Radial boundaries for regions which have 3 ϕ segments per octant (in m)	5.024	5.195	5.372	5.557	5.749	5.949
	6.158	6.375	6.601			

* There are a total of 70 counters per octant in the first layer. Dimensions of the second layer scale according to z positions of layers.




FIG. 7-19. Trigger efficiency of the forward muon scintillation counters as a function of transverse momentum.




FIG. 7-20. Trigger efficiency of the central muon counters for an upgraded system with two layers of scintillation counters.

Light collection and readout

Compared to the central counters, those in the forward region have narrower dimensions and are read out by only one photomultiplier. We propose to extend the mean-timing technique to the forward system by locating the single photomultipliers of counters at opposite ends of the scintillators in the two layers. Muons perpendicularly incident on matched sets of counters will give the sharp time response shown in Fig. 7-18.

The baseline design for the forward counters employs either tapered or twisted-strip plastic light guides for channeling scintillation photons to the photomultipliers. Since the counters in a coincident pair have single light guides at opposite ends, Čerenkov light produced in the light guides cannot generate additional trigger signals as is the case for the central region.

Another, potentially attractive option for detecting photons is to use special photomultipliers with long photocathodes. Such tubes have been developed by V. Rykalin's group at IHEP (Serpukhov). In this scheme, light guides are eliminated entirely since the photocathode dimension is matched to the scintillator cross section. It is even possible to segment the electron multiplier of each tube such that several narrow scintillation counters can be read out by a single multi-channel tube. Such tubes are not available in large quantities at the present time and further study is required.

As part of an ongoing R&D program, we are also investigating the feasibility of an entirely different light collection scheme based on wavelength-shifting fibers. In this configuration, scintillation counters of

dimensions $0.010 \times 0.50 \times 0.93$ m would be read out by ribbons of small-diameter, wave-length shifting fibers surrounding the four narrow edges of each counter. In all likelihood, these shorter counters would be read out with just a single photomultiplier, so that the total number of channels would be the same as in the baseline design described above. Initial tests indicate that this scheme could yield an average of about 20 photoelectrons. Whether it is competitive in terms of performance, reliability, and cost remains to be determined.

Rates in the forward system

Estimates of particle rates in the forward system are highly uncertain for two reasons. First, the rate of decay and prompt muons increases rapidly with decreasing angle making the estimates very sensitive to edge effects at the inner boundary of the forward muon system. Second, backgrounds from low-energy neutrons and photons depend strongly on the precise configuration of the beam pipe and surrounding masks and shields. Very roughly, for $1.5 < |\eta| < 2.5$ and assuming an SSC luminosity of $10^{33} \text{ cm}^{-2}\text{s}^{-1}$, our expectations are:

(a) Muons from pion and kaon decay	400 kHz
(b) Muons from charm and bottom decay	100 kHz
(c) Low-energy (neutrons, etc.)	coincidence rate unknown

We are in the process of calculating the rates expected from low-energy backgrounds, and are performing measurements of neutron rates [8]. We have proposed a system with two layers to provide fast trigger roads and with a radial segmentation that allows the exclusion of inner segments if trigger rates are too high. Requiring coincidence between widely separated counters will suppress non-directional, low-energy backgrounds.

7.3.3. Calibrations

The conceptual design of the calibration system for the muon trigger scintillation counters is still in a preliminary stage. The basic scheme is to generate light pulses at each counter whose amplitude can be varied in a well-defined manner to measure the gain and threshold of each channel. Gains can be adjusted by changing the photomultiplier high voltage to insure that the response of the counters to muons is well above threshold.

Light pulses could be generated at each counter by distributing the output of a central laser via optical fibers. An alternate scheme would best mount individual light-emitting diodes on every counter. Whatever calibration system is chosen, it will confirm thresholds periodically by monitoring counter outputs as a function of light-pulse intensity. The system will also be used in coincidence with other trigger test programs to check the relative timing of various trigger components in the tracking, calorimeter, and muon systems.

7.3.4. Upgrades

If operation at higher luminosities or the appearance of unforeseen backgrounds demands a higher rejection capability from the muon trigger counters, two upgrade paths are available. At large angles, a second layer of scintillation counters could be implemented outside BW3-IW3. Such a second layer would provide triggers on the time scale of the beam-crossing period with a transverse momentum response shown in Fig. 7-20. The response has a strong angular (η) dependence because of the choice of counters with constant widths. Chamber modules in BW3 and IW3 will be designed such that a second layer of scintillation counters can be easily added if required.

The situation in the forward muon system is more uncertain. It is possible that upgrades will be necessary to achieve the desired trigger rejection capability. If this is the case, we propose the addition of the Čerenkov counter system described in Section 7.4.

7.4. Čerenkov counters

7.4.1. Introduction

It may be difficult to operate an efficient trigger system for high-energy muons in the forward direction with only the use of standard scintillation counters, particularly at luminosities above the design value. For this reason, we are proposing, as an upgrade option, the incorporation of multi-cell nitrogen gas Čerenkov counters into the forward trigger [9]. These counters are very efficient in the detection of high-momentum muons ($p > 20 \text{ GeV}/c$) and completely insensitive to the low-momentum particles expected as the main background. We now proceed to describe these detectors.

7.4.2. Čerenkov counter parameters

Table 7-5 gives the parameters of the Čerenkov detectors.

Table 7-5
Design parameters of the gas Čerenkov counter.

Parameter	Value
Gas medium	N ₂
Gas pressure	1 atmosphere
Index of refraction ($\lambda = 3500 \text{ Å}$)	1.000309
Width of counter	2.00 m
Inner radius	2.50 m
Outer radius	5.98 m
Width of Čerenkov medium	1.85 m
Mirror reflectivity	85%
Photomultiplier	5" Burle 8854
Expected number of photoelectrons	17.5
Number of cells per counter	200
Number of mirrors per cell	4
Size of each mirror	$\approx 40 \times 45 \text{ cm}^2$
Focal length of each mirror	1.40 m
Čerenkov light ring radius at photomultiplier	$\approx 3.73 \text{ cm}$
Momentum threshold for muons	4.25 GeV/c

The operation of the counter as a trigger for high-momentum muons is illustrated in Fig. 7-21, where we show the alignment of a particular mirror-photomultiplier axis relative to the beams collision point. As a muon is bent by the magnetic field in the forward toroids, the angle of incidence of the muon relative to this axis varies inversely as the momentum. This angle would then be reflected in the number of photoelectrons produced in the photomultiplier, since the Čerenkov light ring would also move off the photomultiplier face, as indicated in Fig. 7-21.

Tests of a prototype counter in a Fermilab test beam indicate production of 10.2 ± 1.0 photoelectrons per meter, in excellent agreement with the calculated value.




FIG. 7-21. Čerenkov counter optics.

7.4.3. Detector simulation and trigger efficiency.

We have used the pulse height distribution as a function of muon incident angle measured in the Fermilab test, to develop a simulation of the detector performance as a trigger counter and have determined the trigger efficiency as a function of muon momentum. Figure 7-22(a) shows Čerenkov light rings on the face of the photomultiplier for the sum of many 300, 50, and 20 GeV/c muons. This case is for the center mirror in the smallest radius layer (see Fig. 7-23), while Fig. 7-22(b) shows the case for a corner mirror in the largest radius layer. These two examples cover the extreme cases of how the Čerenkov light ring appears at the photomultiplier face. The effect of multiple scattering is minimal compared to the toroid bending power; the multiple scattering is about $1/7$ of the bend angle and this ratio is independent of momentum. The multiple scattering effect is clearly seen in the increasing spread of the 50 and 20 GeV/c light rings when compared with that of the 300 GeV/c ring. The scatter in the 300 GeV/c ring is due mainly to spherical aberrations. From these two figures, we deduce that the detection efficiency versus momentum is very similar, and, thus, that the trigger efficiency versus momentum is the same for all mirrors.

7.4.4. Mechanical structure

The counter in each end is divided into eight sections (octants) located radially and perpendicular to the beam line. Figure 7-23 shows a frontal schematic view of such a unit, illustrating the mirror structure and location of the 25 photomultipliers, each photomultiplier observing four mirror cells. The location of the photomultiplier relative to its 4-mirror cell is located so as to minimize spherical and astigmatic aberrations. The distance from the interaction point to the mirror plane is about 16 m. The depth of each octant is about 2 m and the volume is 3.9 m^3 .

All eight octants attach to a "spider", a welded structural frame made of aluminum I-beams. This frame is shown in Fig. 7-24. We plan to have dowels on the spider for precision mounting and positioning of the octants. The spider and the octants are joined on the downstream side. On the upstream side, the octants will be bracketed together on the inner and outer radius to add to the rigidity of the structure.

The octants are constructed as a frame made of structural angle and welded panels (two sides, top, and bottom). There are 25 photomultipliers and 100 mirrors per octant as shown in Fig. 7-23.

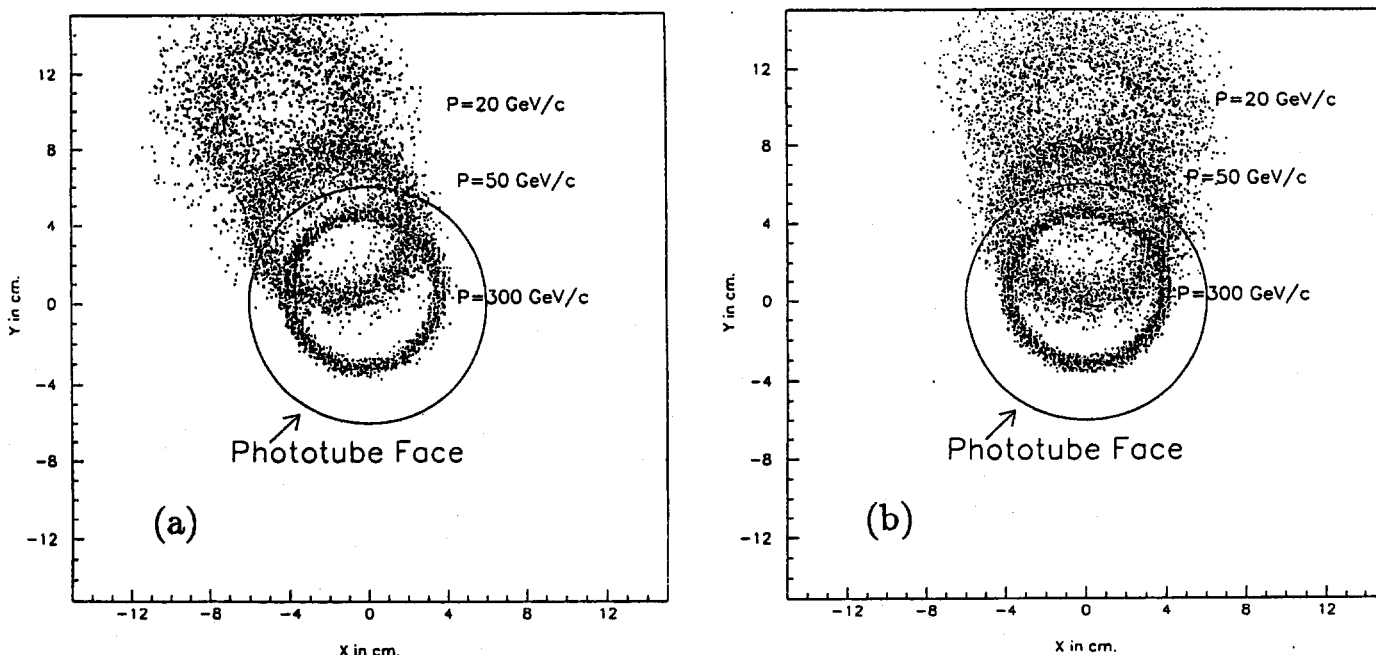


FIG. 7-22. (a) Čerenkov rings produced by the passage of many muons of three different momenta in the center mirror of the innermost layer of cells, and (b) Čerenkov rings produced by the passage of many muons of three different momenta in the edge mirror of the outermost layer of cells.

The counter runs at atmospheric pressure. Assuming an operating pressure of 13 mm of water above atmospheric pressure the net average force on a side is about 230 kg.

7.5. Central chambers

7.5.1. Introduction

The wire chambers in the central region [10] have been optimized to meet the requirements discussed in Section 7.1. Key elements of the design include field shaping to orient the drift direction in each cell to be normal to the path of a high-momentum muon and the arrangement of cells in a projective geometry. Field shaping provides good resolution and two-track pair separation. The use of a projective geometry regularizes and simplifies the trigger electronics and improves the accuracy of the measurements.

The timing of the signal from a drift cell measures the distance from the track to the anode wire. In order to make full use of that information, we must know the location of the wire. There are several steps in getting from the wire coordinates to coordinates related to physics analysis. The first relates the wire segments to the ends of the wire, the second relates the wire ends to reference points on the module containing the wire, the third and fourth steps relate the module reference points to the central tracker reference points and relate these to the beam. The first two steps are specific to the design of the chambers; the last two steps are the subject of the alignment system discussed in Section 7.8.

The basic approach in this design is

1. To separate the fabrication of the drift cells from that of the modules;
2. To make use of the drift tubes as structural elements of the modules; and
3. To establish the alignment and orientation of the drift cells within a module by the drift tube endcaps and the module end plates.

The drift cells are aluminum drift tubes with circular cross section. Each of the cells of a module is a complete and independent, fully functional drift tube with all of the necessary passive interface components built into the endcaps. The proper orientation of each tube, necessary with field shaping, is imposed




FIG. 7-23. Schematic view of one octant of the Čerenkov counter.

during the assembly of the module by mechanical constraints (pins) which are built into the endcap support structure.

The use of round tubes in this design avoids the need for midspan supports for the anode and so also avoids the complications which such supports incur. It allows the drift direction to be optimized for each individual tube location without affecting in any way the external configuration of the body of the tube. The homogeneity of the exteriors of the tubes makes it practical to bond all of the parts of the module together. The tubes then become structural elements of the module. This makes the module rigid without the added weight of stiffening members.

The largest modules are envisaged to be approximately 9 meters by 7.5 meters and to contain the θ , ϕ , and stereo tubes. A cross-sectional view of the corners of two adjacent modules which contain four θ , four ϕ , and two stereo tubes is shown in Fig. 7-25. The gas manifold and gas connections, as well as any electrical leads, are placed in the space between the locating end plate and the protective cover. The anode wire is fixed by the endcap of the tube and located by the CNC-milled endplates.

The most severe rate environment is faced by the BW1 supermodule, since it is not shielded by the barrel toroid. Assuming an SSC luminosity of $10^{33} \text{ cm}^{-2}\text{s}^{-1}$ rate, estimates for this superlayer are [6]

(a) Muons from pion and kaon decay	335 kHz
(b) Muons from charm and bottom decay	105 kHz
(c) Cosmic-ray muons	20 kHz
(d) Particles from showers in the calorimeter	210 kHz




FIG. 7-24. Schematic of the spider frame assembly holding the Čerenkov counter octants.

The total rate from these sources is 670 kHz. Since each drift tube is active for $1 \mu\text{s}$ and there are an average of 1334 drift tubes per layer in BW1, the average occupancy is expected to be around 0.05%. Occupancies outside the toroid are expected to be smaller by a factor of about six.

7.5.2. Construction philosophy

Since we separate the construction of the drift cells from the module construction, the cells can be fabricated at universities and national laboratories, where there is considerable accumulated experience in drift tube construction and operation. Every cell will be tested and certified before it is shipped to the module assembly site. The mounting of the individual cells into the modules will take place at the SSCL.

The precision inherent in this design is determined by the accuracy of a CNC-mill and our ability to maintain the alignment of the modules. The position of the locating holes for one wire is related to the corresponding holes for neighboring wires and to nearby survey holes through common construction on a numerical mill, which has a precision of about $12 \mu\text{m}$. There are several steps relating the wire end location and the local survey targets on the endplate. These include the wire to endcap, endcap to mating surface of the endplate, and the endplate holes to each other and the survey targets. Since each of these will be in the 12 to $25 \mu\text{m}$ range, the overall alignment will be much better than the intrinsic measurement accuracy from the drift time measurements ($250 \mu\text{m}$).






FIG. 7-25. Cross sectional view of the corners of two adjacent BW3 modules containing θ , ϕ , and stereo layers.

7.5.3. Design concept

The basic sensing element is a 90 mm inner diameter aluminum drift tube with 1.6 mm wall thickness. These tubes contain field-shaping electrodes and are arranged in a projective geometry. The tubes in layers 1 and 3 and those in layers 2 and 4 form towers projecting to the vertex. Figure 7-26 shows this geometry for the BW1 θ tubes. The projective tube layers 2 and 4 shown in Figure 7-26 are offset from the 1 and 3 tube layers such that a particle track does not travel through the tube walls in one projective layer and the tube centers in the other. The tube spacing must increase from layer to layer to maintain this projective geometry. The spacing between tubes in a given layer is also increased at larger values of θ , since with close packed tubes, muon tracks pass through two adjacent tubes. Increasing the tube spacing avoids this redundancy and significantly reduces the tube count. This spacing and its increase with θ can be seen in Fig. 7-27. To effectively transfer shear force through a layer, adjacent tubes must be bonded together along a common tangent. When the tubes very nearly touch they are simply glued together; when the separation is greater, spacers are provided. Figure 7-27 gives an overall view of the projective geometry for the θ and stereo layers.

The sense wires within the drift tubes have no support points within the tubes. They are located at their ends by crimped feedthroughs, which are located by CNC-machined holes in the endcaps. The wire tension must be accurately controlled at assembly and the positions of the endcaps must be precisely controlled by the chamber structure. The tension of the wire is measured by pulsing the field-shaping electrodes, which induces a mechanical oscillation (approximately 16 Hz).

All drift tubes are identical except for length. The drift tubes are assembled into modules for each of

FIG. 7-26. BW1 θ drift tubes showing projective geometry.FIG. 7-27. Layout showing increased spacing at large θ angles.

the detector layers and on each face of the toroid. The chamber layout is shown in Figs. 7-28 and 7-29 and consists of three BW1 modules per octant and five BW2/BW3 assemblies per octant. BW1 modules are mounted on rails on the inner surface of the toroid and BW2/BW3 assemblies are mounted on rails on the outer surface. The tube layers are arranged as shown in Table 7-6. The order of ϕ , θ , and stereo layers in any module may easily be changed up to the time when the end plates are machined and the tube lengths are determined. The optimal order of these layers is under study.

Note that the ends of the modules are cut at 22.5° to match the octagonal shape of the toroid and maximize the detector coverage. The breaks between the towers are in planes dividing the BW1 layer into three equal parts. The middle modules on all layers are thus of the same length (7.368 m). The end modules have different lengths as shown in Fig. 7-28. The layout of the θ tubes in IW2 and IW3 is shown

in Fig. 7-30. The ϕ tubes in IW3 are arranged radially as shown in Fig. 7-31. Because this results in overlapping tubes, eight layers of ϕ tubes are needed, arranged as shown in Fig. 7-32. This view at the outer radius of IW3 shows that no extra tubes are required for this configuration. The IW3 stereo tube layout is shown in Fig. 7-33. The interfaces between the barrel system and the intermediate system, and between the intermediate system and the forward system are preliminary and are under study.




FIG. 7-28. Axial view of the barrel region muon modules.






FIG. 7-29. Transverse view of the barrel region muon modules.

FIG. 7-30. Intermediate region θ drift tube layout.FIG. 7-31. Intermediate region ϕ drift tube layout (transverse view).




FIG. 7-32. Intermediate region ϕ drift tube layout (radial view).




FIG. 7-33. Intermediate region stereo drift tube layout (transverse view)

Table 7-6
Layer configuration and channel count.

Modules	Layer	Channels
BW1	4 θ , 4 ϕ ,	10674
BW2	4 θ	7536
BW3	4 θ , 4 ϕ , 2s	26166
IW2	4 θ	1600
IW3	4 θ , 4 ϕ , 2s	11648
Total		57624




FIG. 7-34. Internal structure of a typical module.




FIG. 7-35. Cross section of module edge structure.

The internal structure of a typical module section is shown in Fig. 7-34. The basic principle is that the drift tubes are used as structural members. The outer skins of the modules are made from 6 mm aluminum plate. The first layer of tubes is laid on this plate and fastened to it with a structural epoxy adhesive. The tubes are also glued together at their tangent points. A 1.6 mm aluminum plate is then glued on top of the layer of tubes. Additional layers of tubes are then added as needed for a given module and the final 6 mm outer skin is then expoxied in place. The result is a very strong, light, stiff sandwich structure that can meet the requirement for minimal deflection. The structure around the perimeter of each module has to provide the precision locations for the end caps of the drift tubes, allow for electrical and gas connections and transfer loads from the support points to the module structure. A typical cross section of this structure is shown in Fig. 7-35. The core of this structure is the tube end plate (see Figs. 7-36 and 7-37). This has CNC-machined locations for each endcap plus a number of access holes for electrical and gas connections. At one end of the drift tube, a precision round hole mates with the precise boss on the endcap. Because this end is inserted first at assembly, it has to enter at a small angle. The hole is cut away on each side of the end plate to allow this but still give precise location when the tube is finally oriented (see Fig. 7-38




FIG. 7-36. Typical tube end plate (boss end).




FIG. 7-37. Typical tube end plate (pin end).

for details). The tube is free to rotate in this hole. At the far end, the tube is located by two dowel pins that are pushed into matching holes in the end plate and endcap after the tube is dropped into place. The dowel holes are CNC machined in the end plates and, since they control the orientation of the tubes, they provide the correct orientation of the field-shaping electrodes. The precise location of the locating and dowel holes is vital to the precision of the whole module. Rather than rely on holding the tolerances along the whole length (up to 9.5 m) of the end plate, fiducials will be provided at 2 m intervals. Dowel hole locations will be known with respect to these fiducial marks, which can be surveyed on the alignment stand after module assembly. All machined locating holes and fiducials will be known to CNC mill accuracy ($12\text{ }\mu\text{m}$). Machining will be done in a temperature-controlled environment.




FIG. 7-38. Illustration of (a) the cut-away locating hole which allows insertion of tube endcap boss at an angle while still providing precise location and (b) the cut-away and pins on the flat endcap.

The rest of the peripheral structure of each module consists of top and bottom rails that connect the end plates to the top and bottom 6 mm plates discussed above and to a set of outer plates. The whole arrangement forms a very strong, stiff periphery to the module. The end plates and outer plates are 19 mm thick and the top and bottom rails are 25 mm thick. The top and bottom rails have CNC-machined locations for the end plates and also locate precisely into the outer plates as shown in Figs. 7-39 and 7-40. This edge structure provides the strength and rigidity needed for the module attachment points and for module handling during transport and mounting on the detector. The space shown between the end plates and the outer plates provides room for the electrical and gas connections to the tubes. For the BW2 modules this space is reduced to 10 mm on the module sides which have no tube connections. These modules contain only θ tubes. For the BW1 modules, the outer plates terminate at the top and bottom surfaces of the rails and the whole module is attached to the inner surface of the toroid at points on the top rails.

The BW2/BW3 assembly follows the idea of making full use of the materials in the structure. The two modules must be precisely and rigidly connected together as a unit, which is then installed on the toroid. This is achieved by extending the outer plates of the peripheral structure to connect the modules as shown in Fig. 7-41. Note that the top and bottom rims of the modules connect to precise slots that are CNC-milled




FIG. 7-39. Assembly of the bottom rail to the sloping outer plate and the tube end plate.

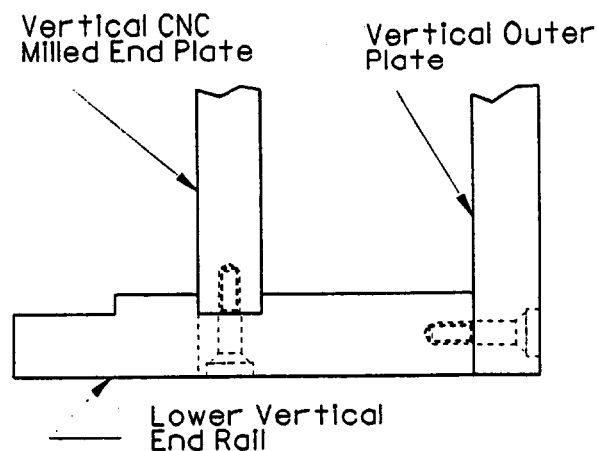


FIG. 7-40. Assembly of the bottom rail to the vertical outer plate and the tube end plate.




FIG. 7-41. Outer plate which connects BW2 and BW3 to form an assembly.




FIG. 7-42. BW2/BW3 assembly.

in the outer plates. As these plates are very stiff in shear and extend around the modules as shown in Fig. 7-42, the result is a very stiff structure which gives accurate alignment between BW2 and BW3.

7.5.4. Structural analysis

The deflections of the structures described above resulting from gravity loads were investigated using ANSYS, a commercial finite element analysis (FEA) package. The details of the methods used are given in Ref. 10. The main goal of the analysis was to confirm the stiffness expected for the structure and to determine if it was possible to mount the modules on true kinematic 3-point supports while keeping all deflections to less than 1 mm. The results were that the maximum deflections are 0.64 mm and 1.0 mm in BW1 and BW2/BW3, respectively, for a three-point support. Thus the BW1 modules can be supported on a true kinematic three-point support. This greatly simplifies alignment of these modules. Also, any reasonable deflection of the toroid will not apply loads to the module but will only move it as a rigid body.

Results for the BW2/3 assembly show that the peripheral structure and the outer plates connecting the

BW2 and BW3 modules form a very rigid structure. Most of the deflection in the assembly consists of the out-of-plane deflection of the plate/tube structure with respect to the periphery and is less than 1 mm at any point. Again these assemblies can be mounted on kinematic supports.

Stresses in the adhesive used to build the tube/plate structure were also studied using a combination of FEA and load tests on sample tube/plate structures. The shear strength of a typical structural epoxy (Ciba-Geigy 1555) was found to be in the range 17.9 to 24.5 MPa. This gives a safety factor of at least 8 for the actual modules under gravity loading.

The other important result of the analysis is that we need adhesive between the tubes as well as between the tubes and plates to maximize the shear stiffness of the modules.

Initial studies of deflections due to thermal gradients show that, while significant, these effects are measurable using a combination of temperature sensors on the modules and deflection predictions using FEA. However more work is needed to quantify the deflection in the thermal environment expected on the detector.

7.5.5. Cell design

The cross-section of the drift tube with field-shaping electrodes is shown in Fig. 7-43. The tube is a thin-wall aluminium extrusion with two T-shaped channels to which the field-shaping electrodes are attached. The field-shaping electrodes are made of coextruded aluminium strip and plastic (Noryl), and are carefully designed to prevent high-voltage breakdown. We choose round tubes because they can be easily rotated to avoid the resolution degradation at large θ , which would occur if the drift field lines were not approximately perpendicular to the detected muon tracks. The upper limit of the tube size is determined by the maximum acceptable drift time and electron diffusion. The design parameters are given in Table 7-7.




FIG. 7-43. Drift tube cross section.

Each measurement station in our chamber system consists of four layers of drift tubes. The design goal for the single tube resolution is 250 μm . In the absence of significant systematic effects, the resolution of each station will be 125 μm . For this resolution to be fully utilized, the wire positions must be alignable to 150 μm . The mechanical design of the drift tube and chamber module are such that this should be possible.

One of the the key features of this design is that the anode wire is only supported at the two endcaps. Tungsten wire of 75 μm (90 μm) diameter tensioned to 750 (1080) grams (75% of breaking strength) will have a sag of 1.2 mm in a 9 m span. The sag is proportional to the length squared, and the length of the

Table 7-7
Drift cell parameters.

Inside radius	45 mm
Wall thickness	1.6 mm
Wire diameter	75 or 90 μm
Wire material	Gold plated tungsten
Wire tension	750 or 1080 g
Electrode width	32 mm
Electrode separation	64 mm
Gas mixture	Ar-CO ₂ (90-10)
Voltage at anode	6.5 kV
Voltage at electrodes	6.5 kV
Gas gain	10 ⁵
Position resolution	250 μm
Double track separation	5 mm

θ tubes varies from 4.5 m to 8.5 m. We have measured the effects of electrostatic forces on a 9 m wire and find that they are well within the stability limits. Corrections for tube sag will be made offline. The trigger is not affected by wire sag as it depends only on time differences between tubes in a projective tower.

The gas mixture is argon-CO₂ (90:10). The drift velocity as a function of electric field is shown in Fig. 7-44. Saturation starts at approximately 800 V/cm and continues to higher fields. The drift field lines and field uniformity are shown in Figs. 7-45 and 7-46. The variation of the drift field within the range 0.8 cm to 4.2 cm is less than 10%. The variation of drift velocity over the entire tube is much less than 1%. The computed x - t correlation curve, given in Figure 7-47, shows very good linearity, and the maximum drift time is 970 ns. Since we operate the drift tube in the saturated mode, degradation of resolution due to environmental factors such as atmospheric pressure variation and temperature change is minimized. The x - t correlation becomes asymmetric if the anode wire is not centered in the drift direction; one cause of this is wire sag. Even in the worse cases, the deviation in drift time from a centered wire is less than 1 ns. The wire sag can also affect the gas gain, but calculations show that this effect should be negligible.

For a given tube diameter and operating voltage, the electrode width and the separation of the two electrodes affects the electric field at the anode wire surface. To ensure that the gas gain not vary more than 10 to 20%, the electrode separation must be controlled to approximately 1 mm. We have measured the electrode separation throughout the length of each of the 30 tubes used in a prototype module and found a maximum variation of 0.5 mm.

Signals from sixteen tubes are transmitted to a connector located on the end plates via single pair twist-to-flat signal cables inside the enclosure of the chamber module. The sixteen pairs of signals are then transmitted to an amplifier/discriminator card via a shielded twist-to-flat cable. A description of the amplifier/discriminator can be found in Section 8.5. We are investigating the possibility of locating the amplifiers directly inside the drift tube endcaps to improve the noise isolation and minimize cross talk.




FIG. 7-44. Drift velocity vs electric field for a 90:10 mixture of Ar-CO₂.




FIG. 7-45. Drift field lines.




FIG. 7-46. Drift electric field vs distance from the wire.




FIG. 7-47. Position vs drift time.

7.5.6. Prototype

We have built a prototype module consisting of thirty 9-meter-long drift tubes (θ cells) and two layers of 50-cm-long ϕ cells. It represents a section of a 9 m by 9 m module. The module was built to test the properties of the design and is currently recording cosmic ray signals. In this prototype, we used 75 mm

diameter (1.6 mm wall) tubes and 50 μm gold-plated tungsten wire. The field-shaping electrodes and endcaps are identical to those described in this document.

To study the performance of the drift tubes and optimize the system design, we conducted several cosmic ray tests. Using a 50-cm long drift tube, the properties of the analog signals were studied. A typical cosmic ray pulse from the amplifier/shaper is shown in Fig. 7-48. The rise time of the signal is 14 ns, and it returns to the base line in less than 75 ns. The gas amplification in this case is approximately 4×10^4 , and the drift velocity 50 $\mu\text{m}/\text{ns}$. Thus, we expect that the two-track resolution will be about 5 mm.




FIG. 7-48. A typical cosmic ray pulse from a prototype 75 mm diameter drift tube (amplifier/shaper output).

Figure 7-49 shows a cosmic ray test setup consisting of four short drift tubes. A Nanometric N277C card is used to read out these drift tubes. Figure 7-50 shows the position-time correlation obtained for one of the tubes by fitting trajectories to the other three. The resolution based on the spread of this distribution is 244 μm .

7.6. The forward chambers

The forward muon detection system has the same goals as the central systems discussed above, but with more emphasis on providing a precise momentum measurement, since solenoidal measurements are weaker in the forward region than they are in the central region.

7.6.1. Special background considerations

Detecting, measuring, and triggering on muons is more difficult in the forward direction than in the central region. The muons tend to have higher momenta and occasionally lose a large fraction of their energy by the catastrophic bremsstrahlung processes. These muons are often accompanied by low-energy electromagnetic debris (electrons and photons) and may confuse the trigger; such effects are under study. To unravel the effect of knock-on electrons, the muon detector must have good multiple track resolution. Calculations with GEANT indicate that the correct hit will be recorded about 95% of the time if the detector has a 5 mm or better two-track resolution.

Charged particle singles rates increase sharply as $|\eta|$ increases. The highest rates in the muon system will be in the tubes at $|\eta| = 2.5$ in FW1, and will be dominated by decay muons. We estimate the rate to be about 20 kHz per tube at design luminosity. Since the maximum drift time in these tubes is 500 ns, this corresponds to an occupancy of 1%.




FIG. 7-49. Cosmic ray telescope showing a typical track.




FIG. 7-50. Position-time correlation measured in the prototype.

Low-energy neutrons may create random hits in the drift tubes as well as in the trigger scintillation counters. Without shielding, the neutron rate in the forward region is estimated to lead to 2 to 3% nominal occupancies in the drift tubes.

The neutrons are produced by low- p_t interactions in the beam pipe, the front and inside surfaces of the forward calorimeter, and the front and inside surfaces of the insertion quadrupole and its protective mask. An iron/polyethylene absorber around the entire beam pipe region protects the forward muon system from the major impact of such neutrons, but no reliable estimate yet exists of the effectiveness of such an absorber for neutrons below 10 MeV. These estimates are in progress.

7.6.2. Special constraints

Access and servicing requirements place special constraints on the forward muon system, which make it more difficult to achieve the 150 μm precision per superlayer that is desired.

The major items that must be considered are:

1. Access to the chamber readout and trigger electronics, and high voltage and gas distribution during short machine shutdowns requires that all these major distribution points be located on the periphery of the detector.
2. Access to the endcap calorimeter photomultipliers and associated electronics, as well as to the inner tracking detector, requires that FW1 be retracted toward FW2 by 1.2 m.
3. Access to FW2 local electronics requires that the supermodule be moved away from the toroids (FT1).
4. Access is required to local electronics for FW4, FS4, the Čerenkov upgrade, FW5, and FS5. The support structure must be designed with this in mind.
5. Access for repairs requires that trapping of modules by the gas distribution and cable runs be minimized.

7.6.3. Mechanical design concept

All of the forward muon superlayers are octagonal in shape. All superlayers measure 4 θ coordinates, but superlayers FW2 and FW5 also measure four stereo coordinates, i.e. two s_1 's ($+7.5^\circ$) and two s_2 's (-7.5°). Table 7-8 lists the z locations, inner and outer radius, measured coordinates, and number of channels for each of the superlayers.

Table 7-8

Forward muon system parameters. The numbers of channels for both ends are given. (Lengths are in mm.)

Superlayer	z_{min}	r_{min}	r_{max}	Coordinate	Tube ID	Channels
FW1	7536	1345	4432	4 θ	42	4390
FW2	9166	1415	5600	2 θ	42	11904
				2 s_1	42	
				2 θ	42	
				2 s_2	42	
FW3 upgrade	11746	1841	5910	4 θ	42	(5788)
FW4	13856	2190	6232	4 θ	57	4310
FW5	18816	3010	8464	2 θ	57	11636
				2 s_1	57	
				2 θ	57	
				2 s_2	57	
Total						32240

The superlayer FW2 is supported from the forward toroid FT1. FW1 is attached to FW2 by a telescoping framework that allows FW1 to move toward FW2 for access to the central systems. The downstream superlayers and scintillation counters FW4, FS4, FW5, and FS5 are supported at their inner radii by a framework cantilevered from the toroid FT2; and, at their outer radii from the nearest major support structure. Additional support trusses are included to maintain the precision and rigidity of the forward system. Since the toroids move upon being energized, any mechanical attachment to them must take this motion into account.

Each of the superlayers is highly modular in design, so that all activities of tube chamber construction, manufacture of chamber mounts, manufacture of support frames, manufacture of local chamber electronics packages, can be carried out independently. The same is true for the testing, measuring and surveying of each of these items. Each set of four or eight coordinate planes will be assembled into a single supermodule, so that superlayers FW1 and FW4 consist of four coordinate planes, while FW2 and FW5 consist of eight such planes. Each of the supermodules consists of eight octant modules that are mounted on an octagonal frame. All local electronics and the distribution centers for gas, signals, and power are completely accessible if one face of supermodules FW1 or FW4 are accessible. In the case of FW2 and FW5 both faces must be accessible to have easy access to all front-end electronics.

The octant modules overlap by being displaced forward or backward in z . Thus there are no dead spaces, either in ϕ or in θ for any of the superlayers. The space created by the overlap is used to house both the local electronics and the utilities distribution boxes.

Figure 7-51 shows an octant module. It consists of a frame holding precision radial mounting plates, similar to those used in the central region. The mounting plates contain holes drilled by a CNC-mill to support the drift tubes. These holes align the drift tube wires to the required precision. The frames will

be made of machined aluminum pieces (or of a single welded piece which is then precision drilled for the location of the mounting holes). About five chamber modules are mounted within each of the octant frames.

The chamber module consists of two doublets of drift tubes and hence consists of about 16×4 tubes, as shown in Fig. 7-52. The second coordinate plane in each doublet is offset by half a tube width in the radial direction. The tubes range in length from 1.2 m to about 7.2 m. The locating precision is achieved by a machined end piece with precise mounting holes. The wire feedthrough holes are precisely drilled with reference to the mounting holes of the muon tube end pieces. Just as in the central region, the field-shaping electrodes are located by dowel pins inserted into precision holes in the drift tube end pieces. The orientation of the electrodes can be chosen so that track always traverses the drift tube almost normal to the field lines. The chamber modules will be made into rigid composite structures by epoxying the tubes together in a fashion similar to what is done for the central region chambers.

The lengths of the tubes do not require extreme tolerance control (± 1 mm) since they are epoxied into the precision end pieces. The wires are strung through the feedthroughs and crimped. Experience shows that the wire location can be maintained to better than $40\text{ }\mu\text{m}$. The end pieces are kinematically mounted onto the radial mounting plates and pinned into place for the correct electrode orientation. The wire location with respect to the octant frame is thus determined by the placement of the wire through the feedthroughs ($40\text{ }\mu\text{m}$), the relative location of the feedthrough hole with respect to the mounting holes of the tube chamber end pieces ($40\text{ }\mu\text{m}$), the relative location of the mounting holes on the radial mounting plate of the frame ($20\text{ }\mu\text{m}$), and finally the location of the radial mounting plates themselves ($40\text{ }\mu\text{m}$).

Finally, the octant modules are mounted onto a support frame, overlapped octant by octant, the even octants being in front of the odd ones. The octagonal support frame is made of extruded aluminum radial spokes and an inner and outer mounting ring to tie all the spokes together. Diagonal extruded aluminum struts furnish additional strength to the entire structure and are located in the space not occupied by the octant modules of chambers. The gas, power, and high voltage distribution boxes are located in this space. Preliminary finite element calculations indicate that such a structure is very rigid. It has a weight less than 35 tons for the largest supermodule, and sags less than 1 mm under full load. Compensation for such a sag can be incorporated in advance into the mounting fixtures of the supermodule. The extra depth in z provided by the octant module overlap contributes to the structural strength of the supermodule.

7.6.4. Muon drift tubes

The muon drift tubes are round tubes with field-shaping electrodes. The baseline design is a smaller diameter version of the central tubes, namely 42 mm inside diameter for FW1, FW2, and 57 mm inside diameter for FW4, FW5, but alternative candidates also exist. All such drift tube candidates will be tested in a high-energy particle beam during the next six months, and their characteristics compared to the baseline design. The drift tube must have a single-track precision better than $250\text{ }\mu\text{m}$ and two-track separation better than 5 mm.

7.6.5. Electronics and gas deployment

The supermodule design integrates the electronic and gas deployment systems within the open spaces provided by the overlap of neighboring octant modules. The local electronics is mounted on the octant module frames. The distribution of gas, high voltage, power, as well as the signal cable switchyard, are located on the periphery of all but the top and bottom octants (see Fig. 7-53). The gas for the drift tubes is brought in by a single pipe that has distribution junctions at each of the supermodules. Each octant in a supermodule can be disconnected by remotely switched solenoid valves. The gas in each octant is distributed from a manifold to sets of sixteen tubes per coordinate plane. Once a gas problem is isolated, the appropriate gas lines can be disconnected, the high voltage turned off, and the rest of the octant can be reactivated.




FIG. 7-51. A forward octant module.




FIG. 7-52. A forward chamber module.




FIG. 7-53. Distribution of gas lines and cables for a forward chamber.

7.6.6. Alignment

Maintaining alignment to the required tolerance is of extreme importance. The critical alignment is that of the radial location of the drift tubes in each supermodule. The radial location must be known to $150\ \mu\text{m}$. The location in z can be good to 1 mm and not affect the trigger or the momentum measurement. Similarly, the plane of a supermodule can be tilted out of its nominal vertical plane by $\pm 1\ \text{mm}$ without causing difficulties. Alignment methods are discussed in Section 7.8.

7.7. Toroidal magnets

7.7.1. Introduction

The barrel toroid is a steel octagonal cylinder with a diameter from outer flat to outer flat of 16.5 m, a thickness of 1.5 m, a length of about 28 m and a weight of 16,400 metric tons. The barrel toroid will be magnetized to 1.8 T. The steel also acts as an absorber to reduce the background for the outer chambers and counters. In addition to its physics functions, the barrel toroid is the structural base for the detector. The diameter and length of the barrel toroid have been driven by the need to provide adequate room for installation and servicing of the detector elements inside the toroid. The thickness is determined by the bend required to give sufficient momentum resolution for the trigger (see Sec. 7.2). Piece size, piece orientation, piece connection, support strategy, and tolerances are selected to provide a safe, and economical design.

The forward toroids FT1 and FT2 are steel octagon-shaped toroids that are located symmetrically at both ends of the barrel toroid. FT1 is 11.2 m across the flats by 1.5 m thick and weighs 1127 metric tons. FT2 is 11.82 m across the flats by 1.5 m thick and weighs 1218 metric tons. The arrangement of the magnets for the muon magnet system is shown in Fig. 7-54.




FIG. 7-54. Muon barrel and forward toroid assembly.

The muon magnet system will use water cooled copper coils to magnetize the iron. A set of sixteen coils with eight turns carrying 4380 A (560 kA-turns) will magnetize the barrel toroid to 1.8 T. There will be four sets of eight coils with eight turns carrying 4500 A (288 kA-turns per set) to magnetize the four forward toroids to 1.8 T. The current will be provided by one power supply for the barrel system and two power supplies for the forward system. Cooling will be provided by chilled LCW distributed through a manifold network on the toroid. A hard-wired system will provide interlocks for coil over-temperature, coil water flow, coil water pressure, power supply local requirements, and ground fault leakage. A monitoring system will provide readings of water temperature, water flow rate, water pressure, current, voltage, magnetic field, and the status of the hard-wired interlocks.

7.7.2. Barrel toroid

The barrel toroid is assembled from 64 short blocks and 128 long blocks. Figure 7-54 shows the final barrel toroid assembly on the supports with the forward toroids and coils in place. Figure 7-55 shows how these pieces fit together into the barrel toroid assembly. The long blocks have a length that tapers from 6.835 m to 9.735 m, a height of 1.5 m, a width of 0.876 m and a weight of 86.0 metric tons. They are used to form the top, bottom, and vertical sides of the octagonal structure. The 45° sides of the octagon are each made from short blocks. These blocks have a length that tapers from 2.733 m to 5.592 m, a height of 1.5 m, a width of 1.752 m, and a weight of 84.4 metric tons. A group of four short blocks and eight long blocks

form one of the sixteen vertical rings that make up the barrel toroid. The blocks within the rings are held together with corner bolts. The rings are connected together by the side plates.




FIG. 7-55. Barrel toroid block assembly.

The size of the blocks is driven largely by a desire to minimize the number of pieces to be handled during installation and the number of surfaces to be machined. The upper limit on size comes from transportation and rigging concerns. The weight of the blocks is limited to the lifting capacity of a 100 short ton crane (89 metric tons). By using the thick short blocks and thin long blocks, both weights come very close to the optimum. The short and long block configuration offers the advantage that all of the corner bolts can be located so that only simple counterbores are required and the axis of the bolt is perpendicular to the joint.

The blocks are made up of trapezoidal plates with a height of 1.5 m and a thickness of 0.147 m. These plates will end up with their large faces perpendicular to the beam axis in their assembled position. The long blocks are made from six plates held together with twelve bolts and two pins. The short blocks are made from twelve plates held together with eight bolts and four pins. There are a total of 768 long plates and 768 short plates. The bolts serve to flatten the plates as well as hold them together. The pins prevent the plates from slipping due to vibration from machining, transportation, and installation. The pins also provide a stronger material for the corner bolts to thread into. There are a total of 2,048 flattening bolts and 512 pins.

The blocks are held to each other with corner bolts and Z plates. There are eight corner joints in each octagonal ring. At each of these joints a long block must be connected to a short block. Each corner interface between a long block and a short block is held in compression by six M56 bolts on the inside and outside of the joint. Each plate has one counterbored-through hole and one tapped hole on each end. The

corner bolts pass through the clearance hole in the plate of one block, perpendicular to the joint and into the threaded hole in the pin of the other block. The bolts used are high strength due to the large moments at the corners. There are some problems with threading high strength bolts into soft magnetic steel. The anti-vibration pin is placed in line with the corner bolts so the threads can be made in the alloy material of the pin. There is also a considerable shear load that must be carried by the joint. The friction at the joint can carry much of the shear but a shear key is included in the design. The long blocks have a 50 mm notch that accepts the 50 mm bump on the short blocks. This shear key also aids in the assembly process. The blocks must be held together in the z direction. This is accomplished with the side plates. Each interface between a block and a side plate is connected with bolts and pins or welding.

The material used for the barrel will most likely be A-87 steel produced in Russia, which is similar to 1010 steel produced in the United States. The bolts should be high strength to reduce the area of steel with poor magnetic properties. Since the pins act as nuts for the high strength bolts, they need to have enough strength so that the internal threads do not fail. Due to their large diameter, the magnetic properties of the pins must be considered in material selection.

The fabrication of the barrel toroid will start with hot rolled steel plates, nominally 0.147 m thick. It may be possible to use side rollers to control the plate width to 1.500 m. The rough shape of the plate would be cut with an Oxy-Fuel torch. It may also be beneficial to rough flame-cut the pin hole. The plates would then have the flattening bolt holes, flattening bolt counterbores, rough pin holes, and possibly a reference surface machined. These pre-drilled plates would be stacked, pressed, and bolted together into blocks. The blocks would be inspected to insure the gaps between the plates are small enough to meet specifications. The plates for the blocks that make up the central two rings will need to be machined on their large faces before assembly to prevent projective gaps at the center of the barrel toroid. The first operation on the assembled blocks will be to bore the pin hole and install a large tapered pin. The pin will keep the relative alignment of the plates before, during, and after machining. The joint faces of the blocks and the large faces of the blocks will need to be machined. Since the short thick blocks span two long block faces at the corner joint, it may be desirable to machine the joint faces of adjacent large blocks in pairs to reduce the step and slope variations that the thick short blocks need to bridge. The other block faces may also be machined depending on the surface quality achieved in the forming and cutting process. After the joint face is machined, the holes for the corner bolts can be made. Machining details for muon chamber track mounting, calorimeter track mounting, coil mounting, end toroid support connections, and barrel toroid support connections will be added as required. The pre-machining of as many details as possible will reduce the assembly work in the hall.

As the blocks are completed, they will be inspected. The first group of blocks produced will be used to do a vertical test assembly. This assembly will be completed and approved before additional blocks are fabricated. The purpose of the vertical ring assembly is to verify the corner joint design and manufacturing. After successful completion of the vertical ring assembly, completed blocks will be checked as ring assemblies horizontally or vertically.

The completed blocks will be transported by boat to a port on the Gulf Coast of Texas. The main storage facility for the blocks will be at the port unless it is shown that major storage at the SSC site is more economical. There will still be a need for some surface layout/storage area for the steel at the IR8 site. Trains will be used to take the blocks in sets of twelve from the port storage area to a side rail near on the site. Heavy haul trucks will be used to take the blocks from the train terminal to the IR site. The blocks will be lowered directly into the hall. A storage area for assembly will be the operating floor area at each end of the detector pit.

There are 192 blocks to assemble into the barrel toroid. At a rate of two blocks per day, the assembly would require 96 work days (less than six months). By starting in the center and working out, it will be possible to use two crews operating the two hall cranes. The first phase of assembly will be to install all the long blocks for the bottom octant. The assembly sequence is discussed in Chapter 13.

The barrel toroid design is being verified through extensive analysis and discussions with a production facility. The analysis plan outlines a process to calculate the deflections, corner bolt loads, plate to plate

shear loads, plate flattening bolt loads, support failure modes, and magnetic map. A general system model including the floor, support, toroid and subsystem loads has been developed using the ANSYS finite element analysis code. The general model uses a coarse fully elastic mesh. The detailed behavior of areas with soft connections is simulated in the general model by locally changing the elastic modulus of some elements in some directions. The selection of which elements to change and what modulus to use is done by comparing a fine mesh non-linear detailed model of the connection region with the equivalent section of the general model. This method has been used for both the corner bolts and the block shear behavior. The general model will be exercised for many different load cases and some design variations. Specific load cases will include various stages of barrel toroid assembly, subsystem installation, lateral loads in z and x , partial failure of the support system, and cycling of the magnetic field. Of course, an additional requirement is that the design be buildable. There have been, and will continue to be, discussions with steel manufacturers to verify that the design is both feasible and cost effective.

Tolerances on the blocks are particularly important with this design because of the large number of interfaces and the three-dimensional nature of the problem. The tolerances come into play, not only by influencing the final shape, but also to determine the loading on many of the bolts. When forces are applied to force the blocks to form closed rings, and to close the gaps between rings, the corner bolt stress levels will increase. The current design has a significant margin on the corner bolt strength to accommodate these loads. Surface location and profile tolerances of 0.2 mm (0.008 in) appear to be achievable with pieces of this size. These tolerances were achieved on the blocks used to construct the ZEUS detector iron. Tolerance stack-up is being studied to understand the maximum deviations expected. The effect on the corner bolt loads was found to be acceptable in a preliminary analysis, but further analysis will be done after the tolerance stack-up is more fully understood.

7.7.3. Barrel coils

Each barrel coil has an overall length of 28.702 m, width of 1.623 m, height of 0.530 m, and weight of 6.84 metric tons. The coils are manufactured from 35 mm \times 51 mm bus bars of CDA 102 copper, with a 19 mm hole. The net current density is 2.92 A/mm². The bus bar is held together and insulated by 0.25 mm B-stage epoxy-impregnated tape half-lapped around the conductors and cured in place. The construction sequence for the barrel toroid requires the coils to be installed after the steel is completed. To accomplish this, each barrel coil is made from two halves with eight conductors that are joined on each end with conductive shims, insulating shims, clamp plates, and bolts. When the two halves are joined, the conductors closest to the steel on one end and the conductors farthest from the steel on the other end form one electrical path, while the remaining conductors form a second electrical path. The DC power is input into each electrical path where it is farthest from the steel. The coils are interconnected electrically into two separate circuits with their own distribution system by bus-bar interconnects at each end of the barrel toroid. Each circuit requires a minimum of 4380 A at 212 V for a power of 929 kW (1,858 kW total). Both ends of each of the sixteen conductors in the joined coil have a water fitting. Four water paths are created by joining groups of four conductors together with hoses to form a water circuit. On each end of the barrel toroid there is one ring shaped supply manifold that delivers 18.3°C LCW and one ring-shaped return manifold that collects 29.4°C LCW. The total LCW requirement will be 0.040 m³/s at a pressure of 0.51 MPa. The routing of the water through the conductors is arranged to set up counterflow cooling so the coil appears to the muon chambers as a fairly uniform thermal surface at 23.9°C. The coils will be supported off the steel by twenty brackets per coil spaced roughly 2.8 m apart. No thermal or magnetic shielding is planned at this time, but is under study. These coils, along with the coils of the forward toroids, are shown in Fig. 7-56.




FIG. 7-56. Coils for the barrel and forward toroids.

7.7.4. Toroid support

The toroid support is the structural device which transfers the load of the barrel toroid and detector components to the foundation below. This structure is shown in Fig. 7-57.

The detector components are very sensitive to differential deflection and rely on the barrel toroid for stability. The barrel toroid must behave as a single mechanical and structural unit under the influence of floor motion or other perturbations with no discontinuous motion. The toroid support jacking system will need to maintain the following adjustment parameters to control and stabilize the muon barrel toroid: uniform vertical adjustment, non-uniform vertical adjustment, tilt adjustment, rotation adjustment, and combinations of these.

Monitoring of the deflections and loads at the jack support locations is essential. Vertical movement of the foundation due to heave or settlement will be of a dishing configuration, that is, maximum movement will be under the foot print of the load and minimum movement at the perimeter of the foundation. Heaving or deflection movement of the foundation will be very slow, taking several years to develop. To maintain the required vertical barrel alignment, the movement of the foundation and the supports, as well as the loads at the jacks, will be continuously monitored by position transducers, such as a liquid leveling system or laser-based motion sensors. Load cells will be placed at each jack and the hydraulic pressure will be monitored when adjusting the height. Deflection under the detector will be constantly monitored by an automatically controlled system. Data will be monitored automatically and alarms will occur if maximum deflection and load differentials are exceeded.

The alignment of the detector over a period of about one year must remain stable in the millimeter range so as not to lose the alignment precision of the tracking components to the beam. Motions smaller than about a millimeter can be corrected in offline analysis.

The configuration of the toroid support consists of two inclined longitudinal plates resting on two continuous plate girders, which are held together laterally with tie plates and supported vertically by hydraulic jacks. Lock nuts on the hydraulic jacks are required to support the load in the static condition. The inclined plates are held together with a space truss, which is also used to stabilize the support, resist lateral loads, and support two layers of muon chambers. Intermediate adjustable tie plates will keep the girders in alignment. The support is designed to facilitate vertical adjustment of the barrel elevation to compensate for the anticipated floor movements due to settlement of the concrete foundation during assembly of the detector, and due to settlement or heave of the foundation during the life of the detector.

In the final position, the toroid support needs to be placed on a tilt to accommodate the required tilt of the barrel toroid which will be placed parallel to the beam. It is estimated that the installation time for the support system is about four months. The installation of the barrel toroid and support are critical path items and the process is described in Chapter 13.

7.7.5. Forward toroids

The forward toroid combination FT1 and FT2 are located symmetrically at both ends of the muon barrel toroid. The position of the forward toroids is shown in Fig. 7-54 and the details of the forward toroids are shown in Fig. 7-58.

Each forward toroid is assembled from sixteen 1.5 m thick blocks weighing about 80 metric tons each. Each block is built from vertical plates bolted together by high strength bolts. The outer plates of the blocks are thicker (about 0.2 m) and the inner plates are thinner to ease flattening, and to minimize gaps between plates. The outside plates have recessed holes for the longitudinal bolts. The two radial surfaces of each block are machined for proper magnetic continuity and fit. The bearing areas of the bottom blocks and the contact areas of the tie plates between FT1 and FT2 are also machined. For uniform flux distribution the eight radial surfaces between the octants are separated by nonmagnetic plates tapering from 26 mm to 2.6 mm (approximately) from the inner to the outer corners of the toroid.

Consideration is also being given to a round, rather than octagonal forward toroid. The design and




FIG. 7-57. Barrel toroid support assembly.




FIG. 7-58. Forward toroid assembly.

assembly of round toroids would be similar to that of octagonal ones with the major difference being that the nonmagnetic tapered plates would not be required for a round toroid.

The forward toroids are assembled separately from each other. When assembly operations are complete the two toroids at each end of the detector are rigidly fastened to make one unit.

The forward toroids are supported on two thick bracket plates resting on steel pads in the barrel toroid. The supports are equipped with skate rollers to allow movement when opening and closing the detector. Continuous track plates will be embedded in the concrete floor and in the barrel toroid to accommodate the rollers. The rollers may be rotated in any direction to move the toroid sideways should it become necessary.

A transfer bridge structure or grillage structure will be used to bridge the space between the barrel toroid and the operation floor for installation of the calorimeter, forward toroid, and other mechanical equipment. This bridge structure will be designed to induce minimal load and stress impact on the barrel toroid.

7.7.6. Forward coils

Each forward coil on FT1 has an overall length of 4.263 m, width of 1.612 m, height of 0.416 m, and weight of 1.51 metric tons. Each forward coil on FT2 has the same overall dimensions except that the length is 4.223 m. The coils are manufactured from the same conductor and tape as the barrel coils. However, there is no need (or room) for a mechanical joint in the forward coils since they are installed as the steel toroids are constructed. The forward coils are formed as a double helix so there are two electrical and two water paths. The DC power is input into each electrical path through leads on the outer radius of the toroids. All the coils on one forward end are interconnected electrically into a single circuit with its own power supplies by bus-bar interconnects around the toroid. Each of the two circuits requires 4500 A at 81 V for a power of 365 kW (729 kW total). Each helix in the coil is a separate water path. LCW is delivered at 18.3°C from the ring-shaped supply manifold on the barrel toroid and collected at 29.4°C by the return manifold. The total LCW requirement will be 0.016 m³/s at a pressure of 0.60 MPa. The routing of the water through the conductors is arranged to set up counterflow cooling so the coil appears to the muon chambers as a fairly uniform thermal surface at 23.9°C. The coils will be supported off the steel by eight brackets per coil spaced roughly 3 m on the sides and 1.3 meters apart on the ends. No thermal or magnetic shielding is planned at this time.

7.7.7. Magnetic field analysis

We have calculated the field map of the barrel toroid and the fringe fields near it.

In this calculation we have used the B vs H curves for 1010 steel. We have also assumed perfect toroids with no cracks, holes, variations in material properties, or other defects. The operating point of the barrel toroid is defined by specifying that the integral of the field along a line through the center of a flat of the toroid equals $1.8 \text{ T} \times 1.5 \text{ m}$.

Because of the natural symmetry of the octagon, it is only necessary to calculate the field for a 1/8 sector. In our model, the coils are placed reasonably near and on either side of the apex of the octagon. The results of these calculations are shown in Fig. 7-59. The contours show equal values of the magnitude of the magnetic field spaced by 0.1 T. The fringe field in contours of 0.01 T intervals are shown in Fig. 7-59b.

7.8. Alignment and position monitoring system

7.8.1. Introduction

The alignment monitoring will be established through a series of hierarchical position measurements to determine the relative position of fiducial marks. The alignment system will be a distributed, redundant sensor system with verified long term stability. We presently have an R&D program in place to test various ideas, competing technologies and sensor systems. Here we present a model alignment monitoring scheme based on recent test results. Further R&D will determine the final system to be applied to the detector.

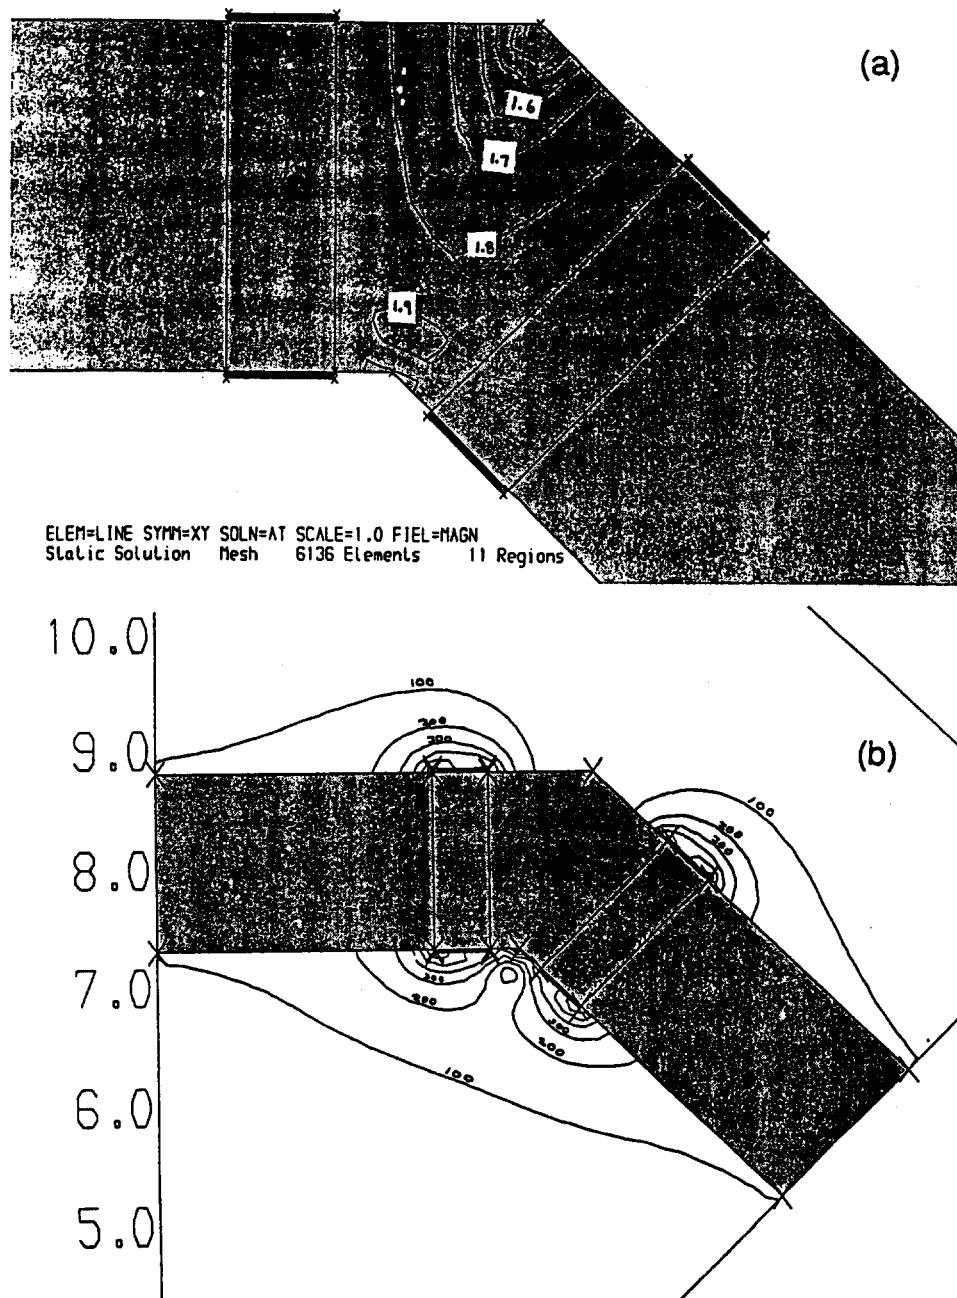


FIG. 7-59. (a) Magnetic field inside the barrel toroid and (b) the fringe field near the coils.

The physical interpretation of the alignment monitoring system will be accomplished through a computer model of the detector in which the relationship between the fiducials and references are modeled with the positions established by the sensor outputs. Temperature measurements coupled with modeled functional dependencies will be used to correct the data.

The toroids play a pivotal role because they support the muon system detectors and separate the inner and outer groups of detectors. For a proper transfer of alignment information, our approach requires the monitoring of the toroid deformation.

The detector coordinate system will be established by transfer from the surface geodetic network. It is estimated that the initial detector axis may be parallel to the accelerator beamline, but maybe shifted by up to 3 mm. With an adjustable base, we propose to roll, elevate, and pitch the barrel toroid to correct

the relative accelerator-detector positions, if required. The goal is to position all of the muon detectors to within 2 mm of the nominal location within the detector coordinate system.

A detector survey installation reference system will be established and maintained by a three dimensional matrix of cross-referenced survey positions in the hall, which are externally referenced by anchored monuments. Secondary stations will be established in the detector system on various stable elements and at appropriate floor locations. A more detailed description of the SDC global survey is available [11].

7.8.2. System support

The base of the barrel toroid must be maintained in its original construction shape for its mechanical integrity and for the stability of detector elements and reference positions. Adjustments will be made to keep the magnet to within 3 mm of nominal position over its length. To achieve this, the magnet base is set on a string of hydraulic jacks with load sensors as shown in Fig. 7-60. The shape and elevation of the base will be monitored by a chain of level detectors along the base that are referenced to floor monuments. This system is capable of point measurements to 50 μm , with a jack positioning tolerance of about 500 μm . Strings of inclinometers on various octants of the magnet will directly monitor its shape to insure mechanical integrity and to provide alignment monitoring system consistency checks. Survey references on outer surfaces and both ends will be established. The inside and outside rails on the magnet that locate and support the barrel muon supermodules will be located relative to the iron to 500 μm . The temperature distribution and shape of the magnet will be monitored.




FIG. 7-60. Liquid level detector system on the barrel toroid base and inclinometers on the iron.

Supports for the intermediate trusses will be mounted and measured with similar tolerances on the ends of the barrel toroid using a transfer fixture with pinned references to the outer barrel toroid rails. The fixture will be used in the surface assembly facility to pre-align and reference the intermediate truss assemblies and establish the relationship of the fixed point support and the IW3/IW2 supermodules.

In the construction of the forward toroids, supports for the forward supermodule rings will be mounted in a plane on all faces of FT1, FT2, and the rear inner shield assembly. The primary fixed support point

for each supermodule layer will be located to an accuracy of $500\ \mu\text{m}$. Again, the magnet temperature distributions will be monitored. Related reference points will be established on the outer boundaries of the toroids. These will be used to establish a set of targets and sensors for radial distance and relative transverse position measurement systems mounted on the BW1 supermodules and in the intermediate trusses on the barrel toroid. Another use of these targets is to measure and align the forward systems relative to the barrel system. A schematic of the model alignment monitoring system is shown in Fig. 7-61.




FIG. 7-61. Schematic illustration of the supermodule position monitors and the monitors that measure relative positions of detector subsystems.

The FW1 supermodule rings will be connected to FW2 and its support from FT1. FW1 will articulate on precision rails, to allow access to the endcap calorimeters, and will relocate on stops. To measure and monitor the position of the FW1 ring relative to FW2, we will measure the radial and z gaps between FW2 inner octants and inner shield ring references projecting from the endcap calorimeters with linear analog proximity sensors. Our laboratory alignment tests indicates a long term resolution of $25\ \mu\text{m}$ on such sensors. Test results are given in a supplementary report [12].

7.8.3. Supermodule referencing

Supermodule bulkheads will include CNC-machined survey reference points and target holes, in addition to mounting holes for proximity position monitoring device assemblies. Surveys of the references on each bulkhead segment and temperature monitoring of the bulkheads will link the positions of all wires in the assembly. The positions of wire cell holes in bulkheads during the machining process should be known to better than 40 μm .

To monitor supermodule positions on the barrel toroid, removable, interlocking two-piece rectangular stainless steel tube structures ("fence posts") will be installed in magnet iron wormholes and extend from the inside of the BW1 layer to the outside of the BW3 layer. These will be installed at the four corners of the BW3, BW2, and BW1 supermodules and can be surveyed during installation at the two post ends. The outer BW3 end will always be exposed for survey. The fence posts will be supported in bushings at the inside and outside face of the magnet, and the joint will be inside the iron. Inside the fence post, two narrow angle (2°) orthogonal inclinometers, with few arc-seconds resolution, will monitor the tilt angles of the post in the flexing magnet. There will also be an optical lens at this point. The two endcaps of the post will contain an IR LED and a continuous x and y transverse position sensing photodiode [13]. Because of the 3 m cantilever outside the magnet (1 m inside), this system will primarily monitor the axis and shape changes of the post on the outside of the magnet. Thermal sensors will define changes in post length and radial ear positions. The assembly is shown in Fig. 7-62. Each of the four supermodules adjoining the post has a corner set of three linear, analog output, inductive, proximity sensors. Two measure the orthogonal transverse gap between the sensors and the walls of the post, and the third measures the radial gap between the sensor and a metal ear on the post. These sensors will be referenced to the wires in a surface facility alignment setup. Laboratory tests over several months demonstrate a readout stability and resolution of 25 μm for these sensors when small temperature corrections are applied [12]. Typical sensors have a linear readout over a 2.0–6.8 mm gap (2V/mm) for a stainless steel post. This range of position tracking is believed to be adequate. However, thin foil skins with plastic backing on the posts can extend the linear gap range to 3.3–8.9 mm [12]. Longer range, nearly linear capacitive sensors are being evaluated for use in high magnetic field regions. Every tower will thus be linked directly to all of its eight surrounding towers by the common fence posts. The fence posts internal monitoring systems and the proximity sensor assemblies for supermodules will be calibrated on a precision optics bench fixture which can be oriented in the different octant orientations.

In the surface facility, each barrel tower (BW1, BW2, and BW3 layers) will be surveyed in its detector equivalent orientation on a rigid open frame structure simulating the magnet iron section. On this frame, the BW1 supermodule and the BW2-BW3 outer tower are supported by their rail fixtures. Additional survey references on the tower system will be established that will be visible during and after installation on the detector. Alignment monitoring devices will be mounted on the corners of supermodules and referenced to the wire positions. Finally, precalibrated alignment monitoring fence posts will be installed on the frame as a mock magnet section. The BW1, BW2, and BW3 relative alignment monitoring devices will be cross checked relative to the fence posts and survey measurements. Temperature monitors will be installed on the tower assembly and temperature excursion effects will be measured.

The surface assembly of intermediate end trusses is essentially the same as the barrel outer tower except that they will be supported from the magnet ends on in-plane kinematic ball and socket mounts. Pre-alignment of supports will be done on the assembly fixture. Thermal sensors will be installed. Again, each outer intermediate tower will be surveyed in its detector equivalent orientation using the steel octant transfer fixture. Visible survey references will be established. IW2-IW3, IW3-BW3, and IW2-FW4 relative alignment monitoring devices will be installed and calibrated. The intermediate to forward system distance and relative transverse position monitors on the intermediate truss will be calibrated by the survey of forward system targets and sensors against sensor readout in a test setup.

The surface assembly of forward system supermodule rings (or half-rings) from octant modules will be done in the vertical plane. The rings will be supported on three in-plane kinematic ball and socket supports near or on the inner perimeter of the rings. After the first two bilayer assemblies, both faces will be surveyed




FIG. 7-62. Schematic illustration of the barrel tower position monitoring fence post.

to establish module wire locations. Temperature sensors will be mounted throughout the rings. Outer and inner perimeter reference points will be established and temperature calibrated. Subsequent bilayers in a supermodule will be surveyed from one face in the assembly. Outer perimeter alignment sensors will be mounted, referenced, and calibrated on mock targets in test fixtures. Along the center line of empty bilayer octants, radial tungsten or invar wires will be installed from the inner perimeter fixed point supports to outer spring tensioned guides shown in Fig. 7-63. We believe that these are necessary because of the thermal loading of the chambers by the forward toroid coils. Near the outer perimeter, targets on the wires and linear analog proximity sensors on the ring will provide radial expansion measurements of the rings. These will be surveyed and referenced in the assembly. The FW4, FW5 supermodule rings will be referenced to support rails on the absorber shield through support bearings on the assemblies. These systems will be referenced to FT2. If the cantilevered rear inner shield requires a permanent support bridge to the magnet support rails, then half rings of FW4 and FW5 will be assembled and referenced.

7.8.4. Inter-regional referencing

In order to connect the intermediate system to BW1 and BW2, we will link the systems through the ends of BW3 and the barrel fence post system. In the intermediate trusses, the IW3 supermodules carry 1-m long edge pillars that project inward in the z direction to the IW2 and BW3 supermodules. The end positions of these pillars have been referenced to the IW3 wires and to external surface survey targets on IW3 in the surface facility pre-alignment. Each IW2/BW3 supermodule adjoining a post (two octants) has a corner set of three linear, analog output, inductive, proximity sensors for each pillar. Two measure the orthogonal transverse gap between the sensors and the walls of the post, and the third measures the radial gap between the sensor and a metal ear on the pillar. These sensors have been referenced to the IW2 and BW3 wires in surface facility pre-alignment. These sensors will link the radial and transverse positions of




FIG. 7-63. End view of the detector illustrating the vertical wire planes used to monitor the intermediate system trusses and BW1 projection sensors to the outer muon tracking chambers.

adjacent IW3 and IW2 towers. Precision inclinometers on the IW3 supermodules will monitor the planes of the IW2-IW3 towers. Capacitive proximity sensors may also be used to monitor the IW2 supermodules relative to the barrel toroid ends. Also the outer surfaces of the eight IW3 supermodules may be monitored using vertical planes (triangles) of constant tension current carrying wires between three visible mounts. Stray field differential flux transformers (bucking coils) mounted on the IW3 supermodules would be used to monitor their z positions relative to the wires as shown in Fig. 7-63. This system is under development. The intermediate towers will be linked to the barrel system towers along the edges of the two end BW3 supermodules in each octant.

To monitor the relative radial positions and orientations of the FW2, FW4, IW2, and FW5 supermodules, we will install z direction fence posts along the outer octant boundaries which will be supported at two ball and socket joints from FT1 and FT2. One joint will be fixed and one will have z compliance. FW4 will be notched for the fence posts and FW5 position sensing will be done on the inside face. The configuration is shown in Fig. 7-61. Like the barrel fence posts, these will have independent orthogonal inclinometers, thermal monitors, and axis and shape-monitoring optical systems. At each octant boundary, the forward supermodules will have corner sets of three analog output proximity sensors. Two will measure the orthogonal transverse gap between the sensors and the walls of the post, and a third will measure the z gap between the sensor and an ear on the post. These sensors will be referenced to the wires of a supermodule in the surface facility assembly. The relative measurements and calibration on the complete

set of fence posts will be done in the final forward system assembly and survey in the hall outside the barrel toroid. Again thermal effects may be measured and calibrated in the open. Precision inclinometers mounted on the faces and some octant edges of FT1 and FT2 will monitor the orientations of both magnets at all times. A similar set of inclinometers may be installed at the rear inner shield support bridge behind the FW5 supermodule.

A link is needed between the forward and intermediate systems and the forward system and the central tracking system. To link the forward system positions to the barrel and intermediate systems, we will monitor the radial distance between surface reflection targets on four octants of FT1 and the inside surfaces of the ends of the corresponding BW1 supermodules. This distance, on the 100 μm scale, can be measured with precision laser distance measuring devices. Our laboratory measurements indicate a temperature corrected measurement stability of 10 μm in optically shielded conditions [12]. A second set of four laser ranging devices will be mounted on the alternate octants of IW2. These will monitor the IW2-FW4 radial distances. The two sets of measurements will locate the forward systems. In all eight of the intermediate trusses, LED-lens systems at the inside edge of the IW2 supermodules will target focused light spots on continuous transverse position sensing photodiode detector targets on the octant faces of FW4. These will be secondary references for the forward supermodules that will be established in open system survey. The source systems will be referenced to the IW2 supermodules in the surface alignment setup and the source-photodiode calibration will be established in the same setup. From these measurements, we will determine the average z position and rotations of the forward systems relative to the barrel systems.

The primary relationship of the interaction region and the muon system will be established from the tracking system. At the present time, the support structure for the tracking system is not well defined. However, we know that it must include references for alignment and must be thermally and mechanically stable. We propose to link the tracking and barrel muon systems by measuring, through the barrel to endcap calorimeter crack, the transverse displacement of retroreflectors mounted on the outer tracker end bulkheads relative to the lines of sight defined by the fence posts as shown in Fig. 7-64.




FIG. 7-64. Fence post with an extension into the inner detectors.

A secondary accelerator reference can be established and monitored by a relative position measuring system between the FW5 supermodules and the outside of the low-beta quadrupoles. Surveyed linear proximity sensors or laser distance measuring devices mounted on flying buttresses projecting from the FW5 supermodules or the inner shield support bridge could establish these relationships.

7.9. Supertower prototype

A full-scale prototype of a BW2/BW3 module will be constructed at the SSCL in the next year. Depending upon the availability of space, a BW1 module may also be constructed and mounted with the BW2/BW3 module to simulate the final arrangement expected in the detector.

The BW2/BW3 prototype will be about $9 \times 7 \text{ m}^2$. A full scintillation counter array (BS2) will also be fabricated and installed. A smaller, movable array will be located above BW3 to provide a cosmic ray trigger.

Drift tubes will be produced and tested at a number of institutions, shipped to the SSCL, and packaged into BW2 or BW3 layers. The number of drift tube production sites for the prototype remains to be determined, but there will be more than one to test production capability at institutions expected to produce drift tubes during the full fabrication.

Scintillators, phototubes, and bases will be fabricated in Russia and the U.S. The supertower prototype will be used as a test bed for a prototype alignment system. The degree to which final alignment concepts can be evaluated is under study.

Production of this prototype at the SSCL will provide valuable experience for later mass production, train the initial module fabrication crew, and allow high-statistics measurement of drift tube and scintillator performance to be made.

7.10. Assembly and installation

7.10.1. Introduction

The assembly building located at IR8 will be used to assemble the central muon supertowers and possibly the forward system. See Chapter 12 for a description of this building. The assembly building and its immediate vicinity will be used to

1. Receive and store drift tubes coming from remote factories;
2. Receive structural components, cables, gas manifolds, and scintillation counters for module construction;
3. Test the drift tubes;
4. Assemble the drift tubes into large supertower modules which are typically $9 \times 7 \times 1 \text{ m}^3$ and weigh approximately 10 tons;
5. Mount scintillator counters on BW2;
6. Assemble the BW2-BW3 modules into supertowers;
7. Check alignment by mounting each finished supertower in its final position on an alignment fixture; and
8. Store the finished supertowers in a temporary storage area.

Table 7-9 gives the numbers of modules and the numbers of tubes that must be assembled for the central region.

Our schedule calls for the central supertower assembly to take about 2.5 years. Hence we must anticipate a monthly production of about three supertowers using approximately 2000 drift tubes. This means that we must fully process a drift tube every five minutes, assuming one-shift operation, or every ten minutes, for two-shift operation.

The forward muon system modules may be assembled in the assembly building and/or in the south installation shaft head house (see Chapter 12 for a description of these buildings). Studies are just beginning of the space required for the forward system assembly and how its schedule overlaps with calorimeter surface assembly in the assembly building.

A detailed study of requirements for the central region assembly has been completed and is summarized below [14].

Table 7-9
Numbers of modules and drift tubes.

Module type	Numbers of modules	Numbers of tubes
BW1	24	10,674
BW2/BW3	40	33,702
IW2/IW3	16	13,248
Total	80	57,624

7.10.2. Receiving and storing

The drift tubes will be made from cylindrical aluminum extrusions of outer diameter 93 mm and will be up to 9.5 m in length. They will arrive as finished, pre-tested, gas-tight units, ready for final tests and assembly into modules.

The tubes will be loaded into shipping pallets with rigging connections for ease in offloading and storing. These pallets will probably have wheels or rollers on the bottom, or possibly air pads, to facilitate movement, not only into and out of the truck but also into the storage area of the building. The tubes will be stored in these pallets until needed, then rolled or craned to the assembly area.

The storage area must be temperature-controlled. We require a 2-month storage buffer for the central tubes, which corresponds to about 4000 tubes. If each pallet contains 20 layers of tubes, the necessary storage area is about 400 m² (using a packing efficiency of 0.5). Of course, since there will be many types of special tubes (bottom octant, cyrostack opening, access holes, etc.) it is likely that additional tubes must be available and, hence, a larger storage area will be needed.

The scintillation counters must also be stored. They are approximately 0.5 m wide by 1.9 m long with photomultiplier tubes mounted on each end. There are about 70 scintillator counters per BW2 supertower so a two month supply is 140.

Storage must be provided for all the necessary fixtures, both those used during assembly and the installation. Mounting rails and other installation hardware will be stored prior to installation on the iron toroid.

7.10.3. Testing and maintenance

All tubes need to be tested after delivery to SSCL. The pallet will have relatively open ends to allow access to the endcaps so the tubes can be tested. The tubes will be tested for gas tightness, electrical continuity, wire tension, and high voltage. It may be advisable to run a sample of tubes on cosmic rays or sources to insure that each remote factory is producing drift tubes of uniform quality.

Although it is anticipated that in many cases it is best to simply discard faulty tubes, some can be revived with little effort (for example, a touch of epoxy on a leaking endcap). We anticipate that the maintenance area needs to be 14 × 10 m².

More important than the nominal area is the access. These tubes are up to 9.5 m long and do not bend easily. Corners must provide sufficient space to swing a tube around them. Extraction of a 9 m tube from the pallet requires an additional 9 m of free space.

In addition to the dedicated repair spaces, it will be necessary to provide space for a machine shop. While this could serve all of the SDC surface assembly needs, it would be primarily set up to produce hardware for the muon system.

7.10.4. Chamber and supertower assembly and alignment

We anticipate setting up four assembly areas on the floor of the building. Fig. 12-4 (Chapter 12) shows how these might be laid out. At each, area the endplates and bottom membrane are located onto the assembly fixture. The drift tubes are then installed and bonded into place. Once a layer is completed, the upper membrane is bonded into place, and the next layer is installed. When the layers are finished, the outer plates are mounted. For BW2 modules the scintillator planes are also installed during this process. Once the mechanical structure is complete the gas and electrical manifolds are installed and tested.

When BW2 and BW3 (or IW2 and IW3) have been finished, they are brought together and connected by their cover plates or trusses. This should be a fairly simple process since each module is stiff enough to be handled as a single unit with the crane. Once the supertower assembly is complete it is moved to the alignment stand for adjustment and measurement.

After the supertower has been mounted on the alignment fixture, it is rotated to the orientation that supertower will have in the detector. With the correct gravitational stresses, the alignment measures specified in Section 7.8 are performed.

7.10.5. Supertower storage and installation

Once the alignment process is completed, the supertower is lifted by a wheeled vehicle, possibly a rubber-tired stratocarrier, and driven to a temporary storage area that is adjacent to the muon assembly building. A strato carrier would have a frame external to the supertower whereas other alternatives, such as air pads or dollies, would have no external frame and might be able to store the supertowers in a smaller area.

The installation schedule presented in Chapter 13 implies that we must plan on providing enough storage for at least 20 supertowers. With a packing efficiency of 0.8 we need 1600 m². This storage does not need a crane, but should be clean and climate-controlled. There should be space for the final check-out of supertowers before installation.

The first step in the installation process is to mount the rails and rail mounts on the iron toroid. We anticipate the muon assembly building will be useful as a staging area. It will also be useful to stage the supertower installation by attaching the installation fixture onto the supertower and insuring it can fit onto the rails. Some simple final tests will be needed before the assembly is lowered into the hall.

References:

1. K. Bramble *et al.*, Solenoidal Detector Collaboration Note SDC-91-135 (1991); and Y. Arai, *et al.*, Solenoidal Detector Collaboration Note SDC-92-169 (1992). Rates have been calculated for the central region and have been doubled here to account for the forward triggers.
2. H. Iwasaki, Solenoidal Detector Collaboration Note SDC-91-124 (1991); Y. Arai, *et al.*, Solenoidal Detector Collaboration Note SDC-92-169 (1992).
3. S. Errede, R. Gardner, and J. Wiss, Solenoidal Detector Collaboration Note SDC-91-156 (1991).
4. Y. Asano, Solenoidal Detector Collaboration Note SDC-91-00033 (1991); and in "Resolution of a Simple Drift Tube," Internal note, January 1992.
5. Y. Teramoto, Appendix 3 of Solenoidal Detector Collaboration Note SDC-92-169 (1992).
6. Y. Sakai, in *Proceedings of the International Workshop on Solenoidal Detectors for the SSC*, KEK report 90-10; D. Green and D. Hedin, Nucl. Instr. and Meth. **A297**, 111 (1990); Y. Arai, *et al.*, Solenoidal Detector Collaboration Note SDC-92-169 (1992); H. Iwasaki, "Updated Hit Rates," internal note, February 1992.
7. Numbers are from Iwasaki, Ref. 6.
8. K. Johns, private communication. Measurements are being carried out at Fermilab.
9. We would like to thank Michael Marx for his comments to us that led us towards the present design.
10. More details can be found in K. Bramble *et al.*, Solenoidal Detector Collaboration Note SDC-91-135 (1991).
11. D. Veal, Global Survey Alignment for the SDC, in preparation.
12. D. P. Eartly and P. T. Johnson, Studies of Prototype SDC Muon Alignment Concepts and Systems, in preparation.
13. J. Govignon in Solenoidal Detector Collaboration Note SDC-92-170 (1992), Section II.F.
14. J. Brogen, *et al.*, Solenoidal Detector Collaboration Note SDC-92-205 (1992).

8. Electronics

8.1. Front-end electronics, triggering and data acquisition

8.1.1. Introduction to front-end electronics, triggering and data acquisition

The electronics for the SDC detector will be a significant extrapolation from that used in all prior large-scale, high-energy physics detectors. This is true for the radiation environment, channel count, trigger decision time, event rate, and data volume. The parameters of the SSC—beam crossing rate, 60 MHz, and luminosity, $10^{33} \text{ cm}^{-2}\text{s}^{-1}$ —dictate an innovative system architecture. On the average, 1.6 physics interactions occur on every beam crossing. The physical size of the detector implies that particles from as many as four beam crossings are simultaneously contained within the detector. In addition, detector system response times (gas drift times, scintillator fluor decay, etc) increase the memory time of the detector to an even greater number of beam crossings. This leads to the need for a deadtime-less electronics system where recognition and collection of signals from detector elements must go on in parallel with the processing and transport of very large amounts of digital data. Furthermore, the separation of the data into sequential events requires precise timing in at least two ways: 1) all channels need to time-stamp their signals so that all the data from a particular beam crossing of interest can be collated, 2) the system clock, the trigger signals, and other nanosecond-precision signals must be accurately distributed throughout the detector, while taking into account particle time-of-flight, detector signal development time, and electronics delays.

The design of the SDC detector results in $\sim 10^7$ signal channels distributed among ~ 10 different types of detectors (wire chambers, scintillators, etc.), each delivering signals of differing character and of differing timing characteristics. A basic design decision was made to do as much electronic signal processing as possible within the physical volume of the SDC detector. For example, data are stored on-detector while two levels of trigger filtering are performed in order to greatly reduce the volume of event data that must be transmitted from the detector. Providing sufficient bandwidth to transmit all the raw event data (10^7 channels at 60 MHz) would be overwhelming. To support this architecture, dedicated data paths provide trigger data every 16 ns to the first level trigger. With a fixed latency, first level trigger decisions are generated every 16 ns. (Note that throughout this section, the SSC crossing frequency is rounded to 60 MHz and the beam crossing period is truncated to 16 ns). A second level hardware trigger further reduces the number of potential triggers before data are moved off the detector. A final software trigger running in commercial processors uses the entire data set to reduce the number of events written to archival storage.

It has been recognized that the on-detector systems could best be implemented by the wide use of custom integrated circuits designed to satisfy the SDC detector operational requirements. The use of bipolar and CMOS technologies permits the channel densities and low costs necessary to make this system practical. This view, supported by a wealth of SSC [1,2,3] and SDC-specific R&D, has led to the inclusion of custom circuits in the silicon and gas microstrip trackers [4], the straw tube tracker [6], the scintillating fiber tracker [7], the calorimeter and shower maximum detectors [8,9], the intermediate tracker [5], and the muon tracker [10]. The trigger system will also make extensive use of programmable logic ICs.

The SSC may one day operate at a luminosity as high as $10^{34} \text{ cm}^{-2}\text{s}^{-1}$. Design consideration has been made for upgrades to front-end storage, trigger processing power, data transport bandwidth and computer power.

8.1.2. Overview

One of the most important tasks in the development of an electronics system for the SDC detector is the choice of a coherent architecture that meets the requirements of all the particle detectors. The system cost will be minimized and the reliability and ease of debugging optimized if an architecture is chosen that is uniformly applied. The basic scheme of the SDC detector electronics is shown in Fig. 8-1. The blocks "Particle Detectors," "Experimenter's Interface," and "Archival storage" aspects are discussed elsewhere. The other four blocks in the center are within the purview of this section:

- **Front-end Electronics** — The front-end electronics processes detector signals, correlates them with particular beam crossings, filters them in accordance with Level 1 and 2 triggers, digitizes the event data so selected, and outputs event data fragments. The front-end electronics also develops Trigger Data information. Specifics of front-end electronics for various detector types are in Sections 8.3 through 8.8.
- **Data Acquisition System** — The Data Acquisition System (DAQ) assembles the digitized event data fragments into complete event-data records, filters the records through a Level 3 inspection. The DAQ system is covered in Section 8.9 and the Ancillary Controls in 8.10.
- **Online Computing System** — Overall control of the front-ends, trigger system, and DAQ resides in the Online Computing System (Chapter 9) and is supported by the Control/Monitoring Network and Ancillary Controls (Section 8.9). Part of this system is the Archival Storage where complete event records are placed on permanent storage media for access by offline analyses.
- **Trigger System** — Every 16 ns the Level 1 Trigger System receives the trigger data from a new beam crossing and performs logical operations on these data to determine if the event should be discarded or preserved for further processing by the Level 2 Trigger System. The Level 2 Trigger System further processes the Level 1 data making use of additional trigger data from the front ends. The Level 2 Trigger System decides whether the crossing should be discarded or the full event data should be sent to the Level 3 farm for analysis. The Trigger System is described in Section 8.2.

The individual data *paths* are:

- **Detector Signals** — These are a complete set of detector signals that have been sampled and stored by the front-end electronics.
- **Digitized Event Fragments** — Each fragment contains a fraction of the digitized data from a single event. The fragments originate from ~ 280 crates and are carried on ~ 400 optical fiber data links.
- **Complete Events** — These are event data records which contain all the data for a particular event plus any additional information generated by the trigger processors.
- **Trigger Data** — Every 16 ns a subset of the total data for a crossing is sent from the front-end electronics to the Level 1 Trigger System on ~ 5000 1Gbit/s fiber optic cables. These data are also combined with an additional subset of the data for the crossing and sent to the Level 2 Trigger System.
- **Levels 1 and 2 Triggers** — These two trigger signals, the precisely timed 60 MHz clock signal and other fast control signals are distributed by the Trigger System to the front-end electronics and the DAQ system over high speed paths through the Trigger Clock and Control System.
- **Control/Monitoring** — These data paths are used to control front-end parameters such as thresholds, RAMs, etc., and to monitor conditions such as voltages, temperatures and errors/overflows.

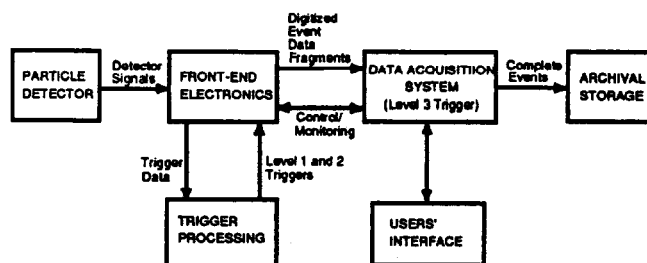


FIG. 8-1. Block diagram of the SDC detector event data gathering and recording electronics system. The arrows indicate principal information paths.

8.1.3. Trigger system

Since the characteristics of the Trigger System influence the architecture of the rest of the system, we describe it first.

The objective of the Trigger System is to reduce the amount of event data that must be handled by making early decisions regarding the utility of the data. This decision is made in three steps—Levels 1, 2 and 3—of successively longer processing times. As shown in Fig. 8-1, the Trigger System receives trigger data from the front-end electronics. Not all detectors contribute trigger data; those that do, and the amounts contributed, are described in Section 8.2. Most of the trigger data have undergone some pre-processing, such as the extraction of particle track segments from a collection of signals from detector elements.

Trigger data arrive at the Level 1 trigger processor every 16 ns on ~5000, 1 Gbit/s fiber optic links. Each link carries a trigger data fragment from a portion of the front-end electronics. Before processing, the Trigger System must sort the data fragments and time align the fragments to create the complete trigger-data set for that event. Since the computations involved in processing are too complex to complete in 16 ns, a pipelined processor is used. In synchronism with the beam clock, a data set is passed through a series of mathematical and logical operations. At the end of the pipeline, a Level 1 trigger decision emerges every 16 ns regarding the physics significance of the data from a beam crossing that occurred approximately 3 μ s before. Another microsecond has been budgeted to propagate signals to and from the trigger system. It is expected that one in 1000 to 10,000 crossings will result in a Level 1 decision to preserve the event from that crossing. Data for the remaining crossings are discarded. The criteria used in this decision are discussed in Section 8.2.

The trigger data for events accepted by Level 1 continue into the Level 2 processor where they are augmented by additional trigger data from other parts of the front-end system. The Level 2 test is more thorough than that of Level 1, and hence takes more time. The time depends on the exact nature of the trigger. For this reason, the Level 2 trigger processor will operate asynchronously. The events are processed in strict time order but the Level 2 decisions occur on average every 10 μ s with the total Level 2 decision time being 10–50 μ s. The Level 2 process is expected to pass only one in 10 to 100 events that have previously passed Level 1. The combined Level 1 and Level 2 triggers reduce the beam interaction rate by approximately 10^{-5} .

Once the decisions are made, the Level 1 and 2 trigger signals must be distributed to the hundreds of front-end electronics crates. This requires an extensive “fan-out” that can control the signal arrival times within a nanosecond or two.

Level 3 decisions are software driven and are executed in commercial processors. They are made on the scale of seconds in the Online Processor Subsystem and use the entire data set of an event. Events that meet the Level 3 criteria are stored on permanent media by the Online Computer System.

8.1.4. Front-end electronics

The term “front-end electronics” is broadly used to denote the electronics that is physically located within the SDC detector. Its basic purpose is to convert sets of signals from detector elements into digital event fragments for those events selected by the Trigger System. It is not possible to devise for all detector systems a uniform implementation of the Level 1 and Level 2 data storage or to dictate the data collection mechanisms before presentation of the data to the DAQ system. However, the same architecture of the two storage levels and their control by the trigger system, as well as a uniform interface to the DAQ have been applied to all front-end systems.

A block diagram showing the functionality of a generic set of front-end electronics is in Fig. 8-2. As drawn, the figure shows a single channel that processes the signals from a single detector element. In practice, large numbers of parallel channels synchronously perform identical processing functions on arrays of detector elements. In the following, the functions of the various blocks are described.

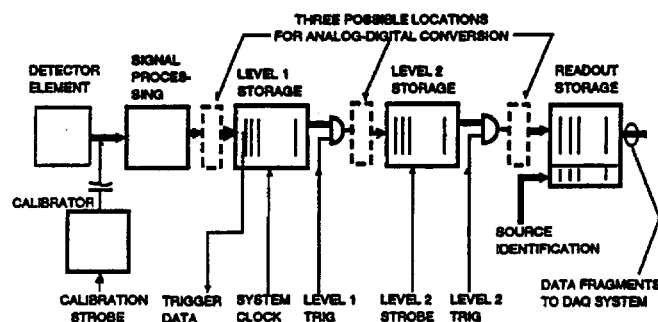


FIG. 8-2. A “generic” block diagram showing the basic functions carried out by the front-end electronics. These functions can be implemented with a number of variations to optimize the performance for a particular detector technology.

The signal processing block typically contains low-noise, high time-resolution amplifiers and signal shapers. Each amplifier receives a signal, almost always an analog signal, from one detector element—wire, Si strip, photomultiplier tube, etc. The signal is “shaped” to extract the best quality time and/or amplitude information from the signal. Digitization or quantization may occur immediately after this stage. The provision for generating calibration signals is typically provided. This feature can be used for channel time and/or amplitude calibration by inserting artificial sets of data to test all or part of the electronics chain.

Whether digital or analog, the signals must be preserved within the front-end system for the $4\ \mu\text{s}$ needed to make a Level 1 trigger decision. In the Level 1 buffer, the correlation of stored data with beam crossing serial number must also be preserved so that the proper data are selected by a Level 1 trigger “accept” decision and so that the data can be correlated with data stored in other parts of the detector.

The very small fraction of data that are selected by Level 1 triggers is entered into the Level 2 buffer. Although the storage time, up to $50\ \mu\text{s}$, is longer than that of Level 1, the volume of buffered data will be much less because of the Level 1 trigger selection. The event data are tagged with salient source information which might include identification of physical origin of the data, beam crossing number, and other items of use in later processing.

A Level 2 trigger “accept” or “reject” decision is received each time the Trigger System makes a Level 2 determination. An accept causes the Level 2 Storage to output the next event data stored within it. This may involve an analog-digital conversion. The data are usually merged with other physically local data and buffered in a queue to be read by the Data Acquisition System. These local data are one of many event data fragments that the DAQ combines to make a complete event description.

Variations in front-end implementation

The front-end electronic system depicted in Fig. 8-2 can be implemented in various ways. Two examples are given.

- **Digitization** — The elements of the data fragments delivered to the DAQ system must be digital. An analog-digital conversion must be located somewhere between the detector element and the DAQ. Figure 8-2 shows the three possible locations. Analog-to-digital conversion prior to Level 1 storage must be performed at the high rate of 60 MHz. The Level 1 and 2 storage are then easily done digitally. This is the choice made for all the tracker systems and the calorimeter Digital Readout scheme. For conversion after a Level 1 trigger accept decision, analog storage must be provided for the duration of Level 1 latency time. Digitization then takes place at the Level 1 accept rate and data storage is digital pending a Level 2 decision. This mode of operation may be used by the shower maximum detector so that full data resolution can be used in the Level 2 trigger. Conversion after a Level 2 trigger accept will be at 1 kHz or less. However, analog storage must be provided for both the Level 1 and Level 2 latency time. This method is used by the calorimeter Switch Capacitor Array readout scheme.
- **Nature of storage** — Figure 8-2 shows two separate storage entities for Level 1 and Level 2. In fact, both storages can be accomplished within a single storage structure as exemplified by the Switched Capacitor Array design described in Section 8.3.2. There, by means of "address manipulation", a single location in an array of storage elements can be assigned the task of Level 1 or Level 2 storage, or can be designated as "available."

8.1.5. Data acquisition system

The Data Acquisition System, described in detail in Section 8.9, contains two major subsystems, the Event Builder and the Online Processor Subsystem.

The Event Builder receives the streams of digitized event data fragments from the front-end electronics over approximately 400 optical fibers links. Each fragment is identified with an event serial number and place of origin. The function of the Event Builder is to correlate and collate the fragments into complete event records. The complete event records are delivered to the Level 3 trigger processor farm.

The function of the Online Processor Subsystem is to implement the Level 3 trigger function in software. A further event reject factor of 10 to 100 occurs here. Those events that are accepted by Level 3 are passed along to the Archival Storage, a part of the Online Computing System (Chapter 9).

The simple block diagram in Fig. 8-1 shows only a few basic data flow paths. Actually, nearly every block will need to communicate in one form or another with every other block. This communication flow is at much lower rates than the event data. To satisfy the need for this type of interchange, a number of slow communication paths are established, the Control/Monitoring Network. Overall management, report logging, and databases for control, configuration, and monitoring reside in the Online Computer System described in Chapter 9. Some uses of the Control/Monitoring Network are: loading front-end operation parameters such as thresholds, data formatting instructions, pedestal suppression values, gains, and trigger criteria; configuring the front ends for calibration data taking; logging error reports; monitoring environmental conditions such as temperature and air-flow.

8.2. Triggering

8.2.1. Introduction

At the nominal SSC design luminosity of $10^{33} \text{ cm}^{-2}\text{s}^{-1}$, an average of 1.6 interactions occur at the beam crossing frequency of 60 MHz. This input rate of 10^8 interactions every second must be reduced by prompt triggers to no more than 1 kHz, a rate that can be analyzed by an online computer farm, and to no more than 100 Hz, a rate which can be written to archival storage, by the full trigger system. The physical size of the SDC detector imposes constraints on signal propagation that combine with electronics technology to require more than $3 \mu\text{s}$ for any primary decision to discard data from a particular beam crossing. Furthermore, the trigger architecture must naturally support an upgrade to handle an increase in SSC luminosity to $10^{34} \text{ cm}^{-2}\text{s}^{-1}$, with the accompanying order-of-magnitude increase in pileup and channel occupancy.

The trigger is the start of the physics event selection process. A decision to retain an event for further consideration has to be made every 16 ns. This decision is based on the event's suitability for inclusion in one of the various data sets to be used for analysis. These data sets include, *inter alia*, di-lepton and multi-lepton data sets for top and Higgs searches, lepton plus missing E_t sets for W , W' , and top physics, inclusive jets for QCD studies, and inclusive electron sets for calorimeter calibrations. Background samples for these data sets need to be made at the same time. In addition, other samples are necessary for measuring efficiencies in event selection. The trigger has to select these samples in real time along with the main data samples.

The measurement of trigger efficiency requires the flexibility to have overlapping triggers so that efficiencies can be measured from the data. The overlaps include different thresholds, relaxed individual criteria, prescaled samples with one criterion missing, and overlapping physics signatures. For example, measurement of the inclusive jet spectrum uses several triggers of successively higher thresholds, with the lower thresholds prescaled by factors that allow a reasonable rate to storage. These triggers overlap in jet energy all the way down to minimum bias events so that the full spectrum can be constructed accurately. The efficiency and bias of each higher threshold can be measured from the data sets of lower threshold.

8.2.2. Trigger principles

The trigger is to be inclusive, local, and measurably efficient, and is to fill the DAQ bandwidth with a high purity stream. The local philosophy of the trigger implies a trigger selection of electrons, photons, muons and jets that relies on local information tied directly to their distinctive signatures, rather than on global topologies. For example, electron showers are small and extremely well defined in the transverse and longitudinal planes. Information from a few calorimeter towers, the shower maximum detector, and a small region of the tracking volume are sufficient for electron identification. The only global entities are neutrinos which are formed from a global sum of missing E_t . For the trigger to be measurably efficient, the ability to measure lepton and jet efficiencies must be built into the trigger architecture from the start. One such tool is overlapping programmable triggers so that multiple triggers with different thresholds and cuts can run in parallel. A second tool is prescaled triggers of lower threshold or weaker criteria that run in parallel with the more strict triggers. A third tool is the prescaling of a particular trigger with one of the criteria removed. The requirement on the use of DAQ bandwidth implies two conditions. First, each level of the trigger attempts to identify leptons and jets as efficiently as possible, while keeping the output bandwidth within requirements. The selected event sample should include all events which would be found by the full offline reconstruction. Hence, the criteria for selecting and defining objects in the trigger must be consistent with those used offline. Second, since the bandwidth to permanent storage media is limited, events must be selected with care at the final trigger level.

We have defined several benchmark processes to judge trigger performance [11]. Not only are these processes interesting in their own right, but they are also typical of final states in more interesting new physics processes. We emphasize that this list of benchmarks is not exhaustive, but does contain representative entries for each type of physics object which the trigger must process. These simple cases are

also the most inclusive triggers, and hence will usually produce the largest rate, making them important for the bandwidth analysis presented later in this section. More complex combinations of the basic physics objects are necessary to produce triggers which are efficient for certain new physics signatures. Examples are di-photons for $H \rightarrow \gamma\gamma$ searches (see Section 3.2.1), the combination of a lepton and di-photon for associated Higgs production searches, and di-leptons for $H \rightarrow ZZ^* \rightarrow 4\ell$ searches (see Section 3.2.2). The benchmarks are:

1. Electrons from inclusive W 's
2. Muons from inclusive W 's
3. Jets at high p_t
4. High p_t photons
5. Missing E_t

The benchmark for processes 1 and 2 is 50% efficiency for W 's which enter the fiducial volume and would not be rejected by offline cuts. The 50% efficiency is almost entirely due to the kinematic cut on lepton p_t . The inclusive lepton triggers will also be efficient for $Z \rightarrow ee$, $Z \rightarrow \mu\mu$. The jet and photon p_t thresholds in processes 3 and 4 must be set such that there is a 1-2 decade overlap in jet and photon p_t with data from experiments at lower \sqrt{s} (for example, CDF). The missing E_t threshold in process 5 is set by having good efficiency for triggering and acceptance for various reaction channels. Examples for missing E_t include the $H \rightarrow 2\ell 2\nu$ mode and searches for SUSY particles.

8.2.3. Trigger structure

In this section, we broadly describe the structure of the three level trigger. The primary objective of the Level 1 is to reduce the rate in a deadtimeless fashion. Level 2 then has the time to search for interesting signatures using a more complete subset of the data than is available at Level 1. Level 2 must reduce the rate so that the event can be readout in its entirety, at which point Level 3 has the full power of the offline code to select and categorize events. The details of the detector information that is used by the trigger are presented in subsequent sections. The data flow diagram is shown in Fig. 8-3.

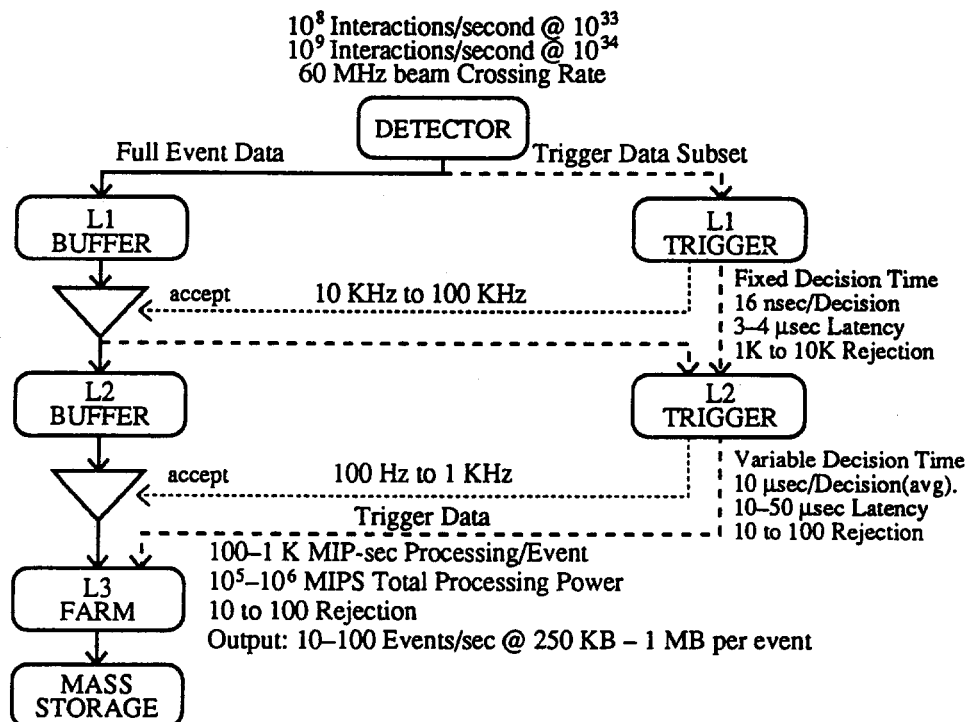


FIG. 8-3. Trigger and Data Acquisition data flow.

Level 1

During the first level trigger decision, all information about the event is preserved in buffers. The first level decision is made, with negligible deadtime, on a subset of the total information available for the event. Made at a fixed time after the interaction has occurred, a first level decision is issued every 16 ns. The Level 1 trigger system must be able to accept a new event every 16 ns. The Level 1 pipeline data storage time is less than 4 μ s [12]. Since signal propagation delays are included in this pipeline time, the Level 1 trigger calculations must be done in many cases in less than 1 μ s. If the first level trigger generates an accept, the event data are moved or assigned to a buffer for the second level decision.

The Level 1 decision is based on the presence of local objects such as photons, electrons, muons, and jets, using information from calorimeters, shower maximum detector sampling, tracking chambers, and muon systems in a given element of $\eta - \phi$ space [13]. It also employs global sums of E_t and missing E_t . Each of these items is tested against several p_t or E_t thresholds. The global compilation of this information is used to decide whether to keep (*i.e.*, trigger on) the data from a particular beam crossing. The Level 1 logic also has the ability to monitor and control trigger rates, hot and dead channels, and other pathological conditions.

Level 2

The second level decision, made on the 10 to 50 μ s time scale, again before the event is read out, also operates on a subset of the total event data. A substantial part of the total data may still be stored in analog form at this point. The second level trigger operates in a series of pipelined steps, each of which is less than 10 μ s, including the transfer of the trigger data from front-end electronics and Level 1 trigger sources. While the total processing time of the second level on an individual event may be as long as 50 μ s, a decision is issued on average every 10 μ s. The combined rejection factor of the first and second levels is 10^5 . The maximum trigger rate of 100 kHz for the Level 1 corresponds to a minimum rejection rate of 10^3 at design luminosity of $10^{33} \text{ cm}^{-2}\text{s}^{-1}$. This maximum Level 1 trigger rate is set by the average time to read information into Level 2 and the average time for completion of processing steps in the Level 2 logic.

Level 2 has access to all the information used in Level 1 since this is stored locally in the Level 1 trigger crates. Consequently, Level 2 can make further combinations and other topological calculations on the digital list of objects transmitted from Level 1. An example is the correlation between muon system and central tracking data. Since the muon trigger information is the last to arrive, processing of this information at Level 1 must be limited to avoid exceeding the trigger pipeline time. The full correlation is done in Level 2 using information stored in the Level 1 system, as well as information from additional chamber planes.

A more important consideration is that detailed information from the muon and tracking systems is not available on the time scale of the Level 1 trigger decision. This information is then used in Level 2. An example is information from the silicon inner tracker, which provides additional track information in both the central and intermediate regions. To limit power dissipation at the detector, the silicon tracker generates trigger information only after a Level 1 Accept.

Level 3

After an event is accepted by the second level trigger all information is digitized and read out. Full event data are used for the third level decision. The event rate into the third level trigger is to be approximately 1 kHz. Level 3, implemented as a processing farm that is designed to achieve a rejection factor of 100, writes ~ 100 events/second to mass storage. The data path to the third level farm is designed to be able to be upgraded to 10,000 events/second. The Level 3 processing does reconstruction and event filtering with the primary goal of making data sets of different signatures on easily accessed media. Data sets of processes with clear signatures (*e.g.*, multi-leptons, W 's, Z 's, high p_t leptons, large missing E_t , *etc.*) will be selected for the offline analyses.

8.2.4. Level 1 trigger function

Level 1 trigger information

Information to be used at Level 1 is sent by the calorimetry, central tracking, shower maximum detector, and muon detector subsystems. Later sections in this chapter describe in detail how the detector subsystems generate this data. We describe below the trigger information which is produced by the baseline front-end electronics, but emphasize the preliminary nature of this. As more detailed simulations are performed and the designs of the contributing detectors and their front-end electronics continue to evolve, we expect that the trigger design will also evolve, and some aspects described below may change. A summary is presented in Table 8-1.

Tracking

The barrel tracking covers the region $|\eta| < 1.8$. The tracking trigger logic puts out two bits of p_t for the highest p_t track for each of $64 \Delta\phi = 0.1$ bins for each half of the barrel. These $\Delta\phi$ bins correspond to those of the calorimeter, shower maximum detector and muon systems. The two bits of p_t allow for two thresholds and the sign of the particle charge. The intermediate tracker covers the range $1.8 < |\eta| < 2.8$ and generates a single bit for a track above a p_t threshold in each of $64 \Delta\phi = 0.1$ bins for each end. The intermediate tracker logic has 4 physical η bins which are combined into a signal η bin per side by Level 1. The fine segmentation of both barrel and intermediate trackers is combined into the 64 bins used in the Level 1 trigger logic. The full granularity of both trackers is used in the Level 2 trigger logic and can be made available to the Level 1 as part of a future upgrade.

Calorimetry

The barrel calorimeter has 32ϕ wedges per end each of which spans 1.4 in η . Four $0.05\Delta\eta \times 0.05\Delta\phi$ EM physical towers are summed to give 0.1×0.1 EM trigger towers. The two longitudinal HAC physical towers are summed into a single 0.1×0.1 trigger tower. There are a total of $64\phi \times 14\eta \times 2$ trigger towers from the barrel. The energy for EM and HAC sections in a trigger tower is digitized separately on an 8-bit nonlinear scale and sent to the Level 1 trigger.

The endcap calorimeter consists of 8 octants per end covering the rapidity range $1.4 < |\eta| < 3.0$. The physical towers are summed to give $64\phi \times 4\eta \times 2$ trigger towers for $1.4 < |\eta| < 1.8$, $64\phi \times 8\eta \times 2$ for $1.8 < |\eta| < 2.6$, and $64\phi \times 2\eta \times 2$ for $2.6 < |\eta| < 3.0$. Each trigger tower has an electromagnetic and a hadronic section. Again, the energy for each section is encoded into an 8-bit nonlinear scale.

The forward calorimeter trigger covers the rapidity range $2.8 < |\eta| < 6.0$. The physical towers for each end are summed to give a total of $8\phi \times 4\eta$ trigger towers with energies digitized on an 8-bit nonlinear scale.

Shower maximum detector

The shower maximum detector has 1024 ϕ strips in $30 \Delta\eta = 0.2$ sections over the range $-3.0 < \eta < 3.0$. There are two barrel sections that cover $|\eta| < 1.4$ and two endcap sections that cover $1.4 < |\eta| < 3.0$. Each section is divided into 32ϕ wedges. In the barrel there are $32\phi \times 8\eta = 256$ bins for each wedge. For the Level 1 there is a single bit for every group of 8 ϕ strips indicating whether any ϕ strip is above a preset "electron" threshold, yielding bins of $0.1\Delta\phi \times 0.2\Delta\eta$. The endcap shower maximum detector is divided into octants. There are $128\phi \times 8\eta = 1024$ bins for each octant. As for the barrel Level 1 trigger, these are combined into bins of $0.1\Delta\phi \times 0.2\Delta\eta$ hit flags. The full granularity of the shower maximum detector is available to the Level 2 trigger logic and can be made available to the Level 1 logic as part of a future upgrade.

Muon detectors

The barrel muon system covers the range $|\eta| < 1.0$. Its scintillators are segmented into $32\phi \times 60\eta \times 2$ ends. For Level 1 the barrel is logically divided into $32\phi \times 5\eta$ regions, $0.2\Delta\phi \times 0.2\Delta\eta$ in size. The muon trigger logic outputs two bits of p_t for the two highest p_t tracks in each region derived from the θ trigger layers in BW2 and BW3. The two bits of p_t allow for two thresholds and the sign of the particle charge.

The intermediate and forward muon systems cover the region $1.0 < |\eta| < 2.5$. Their scintillators are segmented into $32\phi \times 43\eta$ at each end. As for the barrel, the intermediate and forward muon Level 1 trigger systems use information from the two θ layers outside the toroid. For Level 1 this region is logically divided into $32\phi \times 8\eta$ regions, each region being $0.2\Delta\phi \times 0.2\Delta\eta$ in size. Again, two bits of p_t for the two highest p_t tracks in each region are used.

Table 8-1
Level 1 system segmentation and trigger information.

Subsystem	Total channels	L1 segmentation	L1 data
Barrel Track	$64\phi \times 2\eta$	0.1ϕ	2 p_t bits
Int. Track	$64\phi \times 2$ ends	0.1ϕ	1 p_t bit
Barrel Cal.	$64\phi \times 28\eta$	$0.1\phi \times 0.1\eta$	8 bits EM & HAC Energy
Endcap Cal.	$64\phi \times 14\eta \times 2$ ends	$0.1\phi \times 0.1\eta^\dagger$	8 bits EM & HAC Energy
Forward Cal.	$8\phi \times 4\eta \times 2$ ends	$0.8\phi \times 0.8\eta$	8 bits EM & HAC Energy
Shower Max.	$64\phi \times 30\eta$	$0.1\phi \times 0.2\eta$	Hit Flags
Barrel Muon	$32\phi \times 10\eta$	$0.2\phi \times 0.2\eta$	2 p_t bits \times 2 tracks
Int. & Fwd. Muon	$32\phi \times 8\eta \times 2$ ends	$0.2\phi \times 0.2\eta$	2 p_t bits \times 2 tracks

\dagger For $|\eta| > 2.2$, the η and ϕ trigger tower dimensions switch to 0.2 and then to 0.4.

Level 1 trigger rates

In general, the SDC trigger identifies objects such as electrons, photons, jets, muons, and neutrinos and selects events based on interesting combinations of one or more such objects. The trigger uses detector information from local η - ϕ regions to find these objects, other than neutrinos. In some cases information from a single detector component is sufficient while in others the trigger must match local η - ϕ information from several components. The algorithms to identify the various objects are discussed below.

Uncertainties in rates to be expected at SSC as well as detector performance require demonstration that the Level 1 trigger can meet its benchmark performance with an output rate considerably below 100 kHz. In addition, some part of the Level 1 bandwidth must be devoted to the triggers necessary to understand acceptance. We have established a target total Level 1 trigger rate of 30 kHz. There are about eight single trigger-object subtriggers discussed below, such as muons, photons, and electrons, as well as jets and neutrinos. In addition, there are combinations of objects such as di-leptons (ee , $e\mu$, $\mu\mu$) and jets with leptons or missing energy. The rates of these combination triggers are considerably lower than those of the single physics objects. Therefore, we have set a target Level 1 rate of 3 kHz for each of the single objects. We expect one-half the Level 1 bandwidth to be filled by triggers involving the calorimeter system and the other half to be filled by triggers involving the muon system. However, since we do not anticipate an even division of Level 1 bandwidth among subtriggers, we also explore variations in rate and performance. Table 8-2 summarizes the Level 1 rates for single trigger objects and some combinations of them.

In addition to those described below, there will be triggers which do not require the presence of a physics object. Examples of those are minimum bias and sum E_t triggers. Using the resolution and logic

Table 8-2

Representative Level 1 E_t thresholds for triggers at design luminosity, from the simulation of the trigger system described in the text.

Trigger	Threshold
Jet (0.8×0.8 sum)	120 GeV
Electron	20 GeV
Photon	30 GeV
Muon	20 GeV
Missing E_t	80 GeV
Two electrons	10 GeV
Two photons	20 GeV

described below, Level 1 trigger rates have been modelled using 10^5 QCD two-jet events with p_t from 20–200 GeV/c, as well as minimum bias events, at a rate determined by a Poisson distribution whose mean is 1.6 events per crossing. The events were generated by the Monte Carlo program ISAJET; the energy deposits in the calorimeter were simulated by the program SIM3 [14]. The events were integrated over three beam crossings, using the scintillator calorimeter shaping function of the ZEUS calorimeter [16].

Jet trigger rates

Jet triggers require only local calorimeter information. The basic unit of calorimeter Level 1 trigger information is the trigger tower. This is the sum of either the energy in the calorimeter electromagnetic (EM) or hadronic (HAC) compartment in a $0.1\Delta\eta \times 0.1\Delta\phi$ region. Generally, this is composed of the transverse sum of four 0.05×0.05 EM towers or the longitudinal sum of two 0.1×0.1 HAC towers digitized on an 8-bit logarithmic scale.

At Level 1 the trigger identifies jets by summing the digitized energies in the EM and HAC towers in $0.1\Delta\eta \times 0.1\Delta\phi$ to $0.8\Delta\eta \times 0.8\Delta\phi$ regions and comparing these energies with several thresholds. Figure 8-4 shows the trigger rate dependence on sums of energy over threshold for towers of size 0.2×0.2 , 0.4×0.4 and 0.8×0.8 , where the individual $0.1\Delta\eta \times 0.1\Delta\phi$ EM and HAC sums are digitized on an 8-bit logarithmic scale. For a 3 kHz rate, the thresholds are set at 80, 100, and 120 GeV for the 0.2, 0.4, and 0.8 tower sum sizes, respectively. Figure 8-5 shows the efficiency for triggering for these tower sum sizes and thresholds versus the p_t of the jet. This shows that high efficiency for fixed trigger rate is reached at a lower p_t for the 0.8 tower sum than for the 0.4 and 0.2 sums. For the 0.8×0.8 sum, a threshold of 160 (120) GeV yields a rate of 1 (3) kHz and an efficiency of 90% for jets with E_t more than 190 (150) GeV. The final decision on array size awaits a better understanding of the requirements for jet efficiencies at these thresholds and on considerations of the impact of array sizes on the difficulty of implementation.

Electron and photon trigger rates

Electron and photon triggers need information from the calorimeters, shower maximum detectors, and tracking chambers [14,15]. The Level 1 algorithm used to identify electrons and photons involves three items. First, the EM trigger tower sum energy is compared with several thresholds. There are different thresholds for inclusive electrons, di-leptons, photons, and for very high E_t electrons. Second, the ratio of HAC/EM calorimeter energy in the trigger tower sum is required to be small. Third, the energy deposited in the calorimeter trigger towers surrounding this EM cluster may be required to be small if isolation is used in Level 1. In addition, the electron algorithm also requires that a track, with p_t greater than some threshold value, point at the EM cluster and that there is a position match in the ϕ direction between this track and a hit in the shower maximum detector. The cuts applied by these algorithms are relaxed and




FIG. 8-4. The inclusive jet rate versus the trigger E_t threshold for three sizes of sums of $0.1\Delta\eta \times 0.1\Delta\phi$ tower energies digitized on an 8-bit logarithmic scale.




FIG. 8-5. The efficiency for triggering on jets for the threshold yielding a 3 kHz rate versus the actual jet p_t for three sizes of sums of 0.1×0.1 tower energies digitized on an 8-bit logarithmic scale.

finally eliminated for Level 1 triggers with increasing E_t thresholds. The tightest set of cuts is used only to push the electron and photon thresholds to the lowest possible E_t values.

Figure 8-6 shows the expected background rate in the barrel region due to QCD processes plotted against the trigger E_t threshold inferred from a detailed simulation of proposed Level 1 electron trigger hardware [15]. Points are shown for the successive requirements (1) a trigger tower sum with EM E_t over threshold, (2) the tower has a HAC/EM ratio < 0.04 , (3) the energy in the ring of towers surrounding a group of up to 4 towers satisfying (1) and (2) be under a threshold and (4) there be a track with $p_t > 10$ GeV/c pointed at this group of towers.

Figures 8-7 and 8-8 show the efficiency for the electron trigger to select W events and Z events in which an electron enters the barrel region ($|\eta| < 1.5$) [15]. These figures show the efficiency as a function of the rate of background triggers from the barrel region. If the electron E_t threshold is set to 20 GeV, in which case the background rate is 1.5 kHz, the efficiency for detecting such W 's (Z 's) is 79% (82%). If the track requirement is removed, the QCD background rises so that in order to get the same background rate of 1.5 kHz, the electron threshold must be tightened to 30 GeV. This yields 53% W and 64% Z efficiencies.

Since the electron and photon triggers may omit isolation in Level 1, we have also calculated rates without isolation. These results include the calorimeters for both the barrel and endcap regions ($|\eta| < 3.0$). Figure 8-9 shows QCD background rates versus the trigger tower E_t threshold for the inclusive electron and photon triggers. The solid curve is for tower E_t above threshold only; the dashed curve shows the single-photon rate when the condition HAC/EM < 0.05 is satisfied; the dotted curve gives the rate of inclusive electrons with a track of $p_t > 10$ GeV/c matched in ϕ with the calorimeter tower. Figure 8-10 shows the efficiency of the inclusive electron trigger for $W \rightarrow e + \nu$ events. Figure 8-11 shows the QCD background rate versus the trigger-tower E_t threshold for the di-electron and di-photon triggers. The solid curve is for any two towers above E_t threshold; the dashed curve shows the di-photon yield with the requirement HAC/EM < 0.1 ; imposing the further requirement that there be a track with $p_t > 10$ GeV/c matched in ϕ with the tower gives the inclusive di-electron rate shown as the dotted curve.

The combined Level 1 trigger rate of electron, di-electron, photon, and di-photon Level 1 triggers for $|\eta| < 3.0$ is shown in Table 8-3 for various combinations of E_t thresholds on the four triggers. The study indicates that the rate is dominated by the inclusive single electron rates and is well within the 15 kHz target Level 1 calorimeter trigger rate.




FIG. 8-6. Barrel Calorimeter inclusive electron trigger rate vs. E_t trigger tower sum threshold (crosses), with HAC/EM ratio < 0.04 (x's), isolation as described in the text (diamonds) and matched with track with $p_t > 10$ GeV/c (squares).




FIG. 8-7. Barrel calorimeter inclusive electron rate vs. efficiency for detecting W's with an electron from a trigger using an E_t trigger tower sum threshold with HAC/EM ratio < 0.04 and isolation as described in the text (diamonds) and matched with a track with $p_t > 10$ GeV/c (squares).




FIG. 8-8. Barrel calorimeter inclusive electron rate vs. efficiency for detecting Z's with an electron from a trigger using an E_t trigger tower sum threshold with HAC/EM ratio < 0.04 and isolation as described in the text (diamonds) and matched with a track with $p_t > 10$ GeV/c (squares).




FIG. 8-9. Barrel plus Endcap inclusive electron and photon trigger rate vs. trigger tower E_t for tower E_t only (solid), HAC/EM < 0.05 (dash), and HAC/EM plus track with $p_t > 10$ GeV/c matched in ϕ with tower (dotted). The dashed curve represents the inclusive photon trigger rate and the dotted curve the inclusive electron rate for $|\eta| < 3.0$.

Our studies show that conversion electron pairs are the main contributors to the Level 1 trigger rate [14]. Therefore, using the shower maximum detector to reduce the background from the ϕ overlap between a photon and charged track should only slightly reduce the Level 1 rate. Figure 8-12 shows the effect of using the full 1024 ϕ granularity of the shower maximum detector to reduce the background from




FIG. 8-10. The $W \rightarrow e + \nu$ Level 1 trigger efficiency for $|\eta| < 3.0$ from the inclusive electron trigger vs. trigger tower E_t for tower E_t only (solid), and HAC/EM < 0.05 plus track with $p_t > 10$ GeV/c matched in ϕ with tower (dashed).




FIG. 8-11. Barrel plus Endcap di-electron and di-photon trigger rate vs. trigger tower E_t for tower E_t only (solid), HAC/EM < 0.1 (dash), and HAC/EM plus track with $p_t > 10$ GeV/c matched in ϕ with tower (dotted). The dashed curve represents the non-isolated di-photon trigger rate and the dotted curve the non-isolated di-electron rate for $|\eta| < 3.0$.

Table 8-3

Combined Level 1 trigger rate for the main electron/photon triggers in the Barrel plus Endcap calorimeters ($|\eta| < 3.0$) versus various combinations of threshold energies. Where $e \equiv$ electron requires 1 tower with HAC/EM < 0.05 and a track with $p_t > 10$ GeV/c matched in ϕ with the tower, $2e \equiv$ di-electron requires 2 towers with HAC/EM < 0.1 and a track with $p_t > 10$ GeV/c, $\gamma \equiv$ photon requires 1 tower with HAC/EM < 0.05 , and $2\gamma \equiv$ di-photon requires 2 towers with HAC/EM < 0.05 .

Trigger threshold (GeV)				Rate (kHz) @ $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$
e	$2e$	γ	2γ	
20	10	30	20	9.8
20	10	40	20	8.1
25	10	40	20	5.0
25	15	40	30	4.7
30	20	45	30	3.8
20	—	—	—	7.0
—	10	—	—	0.3

overlaps. The figure shows the QCD rate versus the trigger tower E_t threshold for the inclusive electron trigger. The solid curve is for tower E_t above threshold only; the dashed curve shows the electron trigger rate expected with the requirements HAC/EM < 0.05 and a track with $p_t > 10$ GeV/c matched in ϕ with the tower; finally, the dotted curve results from requiring the track to match a shower maximum detector ϕ strip. Since the Level 1 inclusive electron E_t is reduced only slightly, the use of the shower maximum detector is restricted in the design to confirmation of the calorimeter electromagnetic energy. For this




FIG. 8-12. The effect on the Level 1 rate of requiring a shower maximum detector ϕ strip over a threshold to match to a track with $p_t > 10$ GeV/c at a ϕ granularity of 1024. Shown are the rates vs. trigger tower E_t for tower E_t only (solid), HAC/EM < 0.05 plus track with $p_t > 10$ GeV/c matched in ϕ with tower (dashed), and HAC/EM plus shower maximum detector-track match in 1024 ϕ bins (dotted).




FIG. 8-13. The QCD Background trigger rate versus the missing E_t threshold for the simulation of the calorimeter trigger system described in the text.

reason the shower maximum detector is read out in 64 ϕ bins in Level 1.

Muon trigger rates

The Level 1 muon trigger will come from signals from the muon scintillator and muon tracking chambers located outside the muon toroid. The scintillators provide the timing tag needed to associate a muon with a specific crossing.

The muon trigger relies on a scintillator signal in a coincidence confirmed by a projective chamber wire pair. The trigger granularity is derived from the scintillator segmentation. There are 32 $\Delta\phi = 0.2$ and 5 $\Delta\eta = 0.2$ bins for each half of the barrel region ($|\eta| < 1.0$), 2 $\Delta\eta = 0.2$ bins in the intermediate region ($1.0 < |\eta| < 1.4$), and 6 $\Delta\eta = 0.2$ bins in the forward muon region ($1.4 < |\eta| < 2.6$).

The dominant sources of signals from the muon system are low p_t muons from light quark decay. As is discussed in Chapter 7, the muon p_t threshold curve folded with the anticipated muon spectrum predicts a trigger rate between 3 and 6 kHz for a p_t cut that is 90% efficient at 20 GeV/c. The sharpness of the muon p_t threshold curve is limited by multiple scattering and the spread in the interaction point along the z direction. In order to trigger on muons with lower p_t than would be possible using muon system information alone, we can sharpen the p_t threshold by matching the muon candidate on a scale of $\Delta\phi = 0.2$ with a track in the outer tracker. For muon candidates at lower p_t , we can also require the absence of a associated hadron shower with a calorimeter-based isolation cut.

Missing E_t trigger rates

Neutrino identification consists of calculating the event missing E_t vector and testing it against a threshold. The calorimeter trigger calculates both sums of E_t and missing E_t . Only trigger towers above a certain threshold are used in these calculations. The transverse energy vector components are calculated from the 8-bit logarithmic-scale digitized HAC and EM pulse heights converted to a linear 12-bit scale, and multiplied by entries in lookup tables containing the tower angular coordinates. The HAC and EM sums are then combined into single tower sums. Tower sums over the threshold are routed into digital summing networks. Figure 8-13 shows the dependence of the QCD background trigger rate on the threshold in missing E_t . Missing E_t thresholds of 80 and 100 GeV give trigger rates of 3 and 1 kHz respectively.

Level 1 trigger logic

The Level 1 trigger logic operates on data carried on optical fibers from the front-end electronics on the detector. These data are in the form of basic information about physical variables, such as track segments over a certain p_t or energy in the calorimeter towers, or hits in the shower maximum detector strips. The calorimeter front-end electronics transmits eight bits of hadronic and electromagnetic energy for each $0.1\Delta\eta \times 0.1\Delta\phi$ trigger tower. The shower maximum detector transmits hits over threshold for $0.1\Delta\phi \times 0.2\Delta\eta$ bins. The tracking detector transmits two p_t bits for tracks in $\Delta\phi = 0.1$ bins. The muon system transmits two p_t bits for tracks in $0.2\Delta\phi \times 0.2\Delta\eta$ bins. These data are used to form conclusions about the presence of trigger objects in the Level 1 trigger logic. In this and the subsequent sections, we describe a specific hardware implementation [18] of the trigger design. This particular implementation serves as our baseline design and is used in estimating costs. Efforts to refine and optimize the design will continue.

The general strategy of this Level 1 trigger baseline design is that only minimal processing be done on the detector. Trigger information is transmitted to the Level 1 regional and global processing logic located in the counting house on the surface. This reduces power, cooling and space requirements for the trigger logic on the detector and allows access to the circuitry. Reliability is enhanced as most of the trigger electronics can be accessed at any time for repair. Commissioning and diagnosis of the trigger system is eased as the basic signals and operations can be observed directly during accelerator operation, when the detector hall cannot be occupied. With most trigger logic in the counting house the detector and trigger geometry are decoupled. Mapping the geometry of the detector onto racks in the counting house enables the adjacent placement of adjacent trigger towers electronics which might be physically separated on the detector (*e.g.*, barrel and intermediate tracking, calorimetry and muon systems). The location of the trigger logic off the detector also helps to decouple the detector mechanical and trigger hardware designs and allows the trigger logic to be mounted in standard racks and crates.

The Level 1 trigger finds electrons by searching for electromagnetic energy in the calorimeter with a small HAC/EM ratio. In order to distinguish electrons from photons, a correlation is made with the outer tracker. The shower maximum detector information is correlated with that of the tracker and/or calorimeter tower to reduce phototube-noise-induced fake electrons. Photons are selected without the track match, but with a shower maximum detector match and a higher energy threshold [11,15]. Higher E_t electrons and photons are selected with reduced or eliminated requirements on the HAC/EM ratio and track matching.

The Level 1 trigger finds jets by comparing the energy in either 0.1×0.1 regions [19] or 0.8×0.8 regions [18] against a series of thresholds. It finds muons by finding muon tracks over two muon p_t thresholds. For lower p_t muons a simple correlation is made between the tracks from the muon system and tracks in the outer tracker in bins of $\Delta\phi = 0.2$. This is used to sharpen the p_t threshold. Lower energy muons are also checked for correlation with relatively quiet 0.2×0.2 regions in the calorimeter. These last two steps are taken to accept muons that would otherwise not trigger.

A baseline design of the Level 1 decision logic is described in detail in Refs. 17 and 18. In this specific realization, the Calorimeter Summation Crates find electrons by searching for towers with a small ratio of hadronic to electromagnetic compartment energy surrounded by quiet towers. They also sum up energy in

0.4×0.4 regions to compose two overlapping arrays of 0.8×0.8 towers. These crates test the energies in the 0.8×0.8 towers against a series of 8 thresholds, and sum the total E_t , E_x , and E_y energies found in the sums of 0.4×0.4 towers. The E_t , E_x , and E_y sums, as well as the number of 0.8×0.8 towers over each threshold are sent to the calorimeter Energy Sum Jet/Threshold Crates. The electron information is transmitted to Electron Match Crates which combine information on the boundaries of regions covered by the Calorimeter Summation Crates to find electrons on region edges. They also check the electron energy against 6 thresholds.

The shower maximum detector and tracking information is sent to the Track/Shower-Max Match Crates where they are matched on a $\Delta\phi = 0.1$ scale. The η coordinate is found from the shower max hit in 0.2 bins. The p_t value is taken from the track p_t bits. The crates output a series of shower maximum detector hits, some of which are matched with tracks carrying a p_t value, and all of which are on a $\Delta\phi = 0.1$ and $\Delta\eta = 0.2$ scale. These hits are sent to the Electron Match Crates. These crates match the shower maximum detector hits with the calorimeter electrons. The shower maximum detector hits are matched in η - ϕ with the calorimeter towers. The track p_t and calorimeter tower energy are compared for hits which have tracks attached. The result is a series of matched calorimeter to shower maximum detector electromagnetic hits without track matches, "photons", and a second series of calorimeter-shower maximum detector hits with track matches, "electrons." Failures to find a shower maximum detector hit to match calorimeter electrons are also reported.

The calorimeter Energy-Sum/Jet-Threshold Crates generate the number of 0.4×0.4 towers over various thresholds as well as the magnitude and components of the total E_t vector. The Level 1 Decision Crate then applies cuts on the energy sums and towers over threshold to form the trigger. The Electron Match Crates forward the number of "electrons" and "photons" over various thresholds as described above. The Level 1 Decision crate then counts the total number over the thresholds to form the trigger.

The muon system transmits the two highest p_t segments over two p_t thresholds to the Muon Level 1 Crates. These crates count the total number of muons over p_t threshold and forward this information to the Level 1 decision crate. The Track/Shower Max Match Crates also ship the highest p_t track from the outer tracker in each $\Delta\phi = 0.1$ bin with p_t threshold and sign bits to the Muon Level 1 Crates [20]. The Muon Level 1 crates count up the number of tracker-correlated muons and forward this information to the Level 1 Decision Crate. The Calorimeter Summation Crates also provide information on whether there was significantly more energy than several times minimum-ionizing in each $0.2\Delta\eta \times 0.2\Delta\phi$ region to the Level 1 Muon Crates. They count up the number of muon tracks correlated with "quiet" calorimeter regions and forward this information to the Level 1 Decision Crate. The Level 1 Decision Crate cuts on the number of uncorrelated muons, tracker-correlated muons, and "quiet" calorimeter-correlated muons over p_t thresholds to form muon triggers. The final trigger decision is converted to a Level 1 Accept by the Global Clock/Control crate, which transmits the decision for fanout to the Local Clock/Control crates and to the DAQ system.

System layout

The Level 1 trigger is set up in a feed-forward structure. Information is not exchanged between crates at the same level in the trigger decision tree. Instead summary information is sent forward to the next level, which combines the information from the crates in the previous level. The entire structure resembles a pyramid, with the front-end electronics at the bottom, followed by regional processing, global processing, and the final decision crate at the top. Each succeeding layer is smaller and represents a condensation of the information from the preceeding layer. Information from different subsystems is only brought together when required for a trigger decision. This architecture minimizes the volume of data and cables interconnecting the Level 1 trigger logic crates and reduces the problems of synchronization and latency that these interconnects engender.

The Level 1 trigger is also designed to enable straightforward diagnosis, analysis of efficiency, and understanding of functionality. Specifically, this means that we avoid major pieces of inaccessible logic with many different inputs that produce a simple yes or no. Wherever possible, the more difficult to access

front-end electronics produces the basic data (trigger "primitives") used in the trigger decision which are transmitted to the accessible trigger electronics where the logical operations are performed. This reduces the debugging of the front-end trigger electronics to verification of trigger primitive generation. The remainder of debugging and diagnosis is then done in easily accessible electronics that can be checked with test trigger primitive generation.

The crate counts and interconnections for the Level 1 trigger logic are shown in Fig. 8-14. The Front-end crates on the detector transmit their trigger information to the Level 1 trigger system off the detector on 1 Gbit/s optical fibers [21]. Every 16 ns, each fiber delivers 16 bits of data. The same type of optical fibers carry the clock and control communication from the trigger clock and control systems to the front-end electronics subsystems. The Level 1 trigger crates transmit data among themselves on short runs of twist-and-flat pair cable. Each Level 1 trigger crate contains a Level 2 trigger interface card that gathers Level 1 information buffered in the Level 1 trigger input cards and transmits this information to the Level 2 after a Level 1 Accept. Each Level 1 trigger crate also contains a processor board, of the standard variety provided by the DAQ group. Its function is to download trigger tables, monitor trigger data and provide a data path to the event builder for trigger information to be supplied with the event record to be written to storage. Each Level 1 trigger crate also contains a Level 1 Clock/Control Board (L1CCB) that provides the 16 ns clock and the other control functions of the Level 1 trigger system.

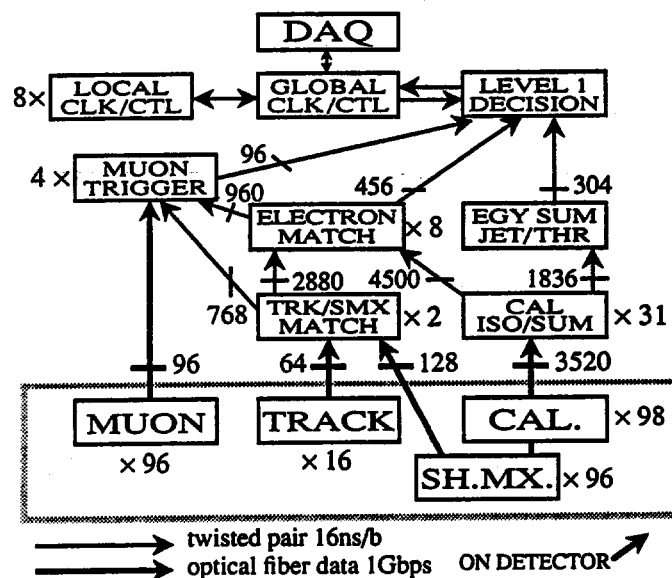


FIG. 8-14. Crates and data links in the Level 1 trigger system. The boxes represent the physical crates and their number while the arrows represent the data links.

Level 1 trigger crates

The functional descriptions of the crates in this Level 1 system are presented below. More detailed descriptions of the individual boards in each of the crates, as well as the crates themselves are contained in Ref. 18. Since the following material is highly technical, the casual reader might want to proceed to the Level 2 trigger discussion, Section 8.1.5.

Calorimeter Summation Crates

These crates find $0.1\Delta\eta \times 0.1\Delta\phi$ calorimeter towers where $E_{EM} \gg E_{HAC}$ which, optimally, are surrounded by quiet towers. These requirements on E_{HAC} and quiet towers are relaxed and finally removed for higher E_t . The electron candidates are reported out on an array of 0.2×0.2 towers. These crates also sum up the total E_t in 0.4×0.4 towers, as well as the total E_t , E_x and E_y in the crate.

There are 31 of these crates: 30 for the barrel and endcap and one for the forward calorimeter. Each crate nominally covers $0.8\Delta\eta \times 1.6\Delta\phi$. It contains 8 Memory Lookup Boards (MLU) that each receive 16 optical fibers that carry eight bits each of hadronic and electromagnetic energy in a 0.1×0.1 tower from the front-end calorimeter crates. Each crate also contains 2 Pattern Recognition ASIC boards (PRA), which receive the information from 4 MLU's, corresponding to a region 0.8×0.8 . These boards perform the search for electron candidates and transmit their results to the the Electron Match Crate. The design details of the MLU and PRA cards, along with simulation results on the performance of this electron pattern logic are found in Ref. 15.

Each crate also contains a double-width Energy Card that covers a $0.8\Delta\eta \times 1.6\Delta\phi$ region. It receives the energy sums in 0.4×0.4 regions from 8 MLU's on a 12-bit scale. In addition to computing the total E_t in 0.8×0.8 regions, it also converts the 0.4×0.4 E_t sums into E_x and E_y sums and adds these up to produce the total E_x and E_y in each 0.8×0.8 region as well. To evaluate the feasibility of making large numbers of pipelined sums of 12-bit energies, we have designed a CMOS 4×12 -bit pipeline adder ASIC that accepts a new input every 14 ns and produces a sum every 14 ns, with a 28 ns latency [22]. The Energy Card also makes combinations of the E_t in each 0.8×0.8 tower to test against thresholds [23]. These results are reported to the Energy Sum Jet/Threshold Crate.

Track/Shower-Max Match Crates

These crates correlate the outer tracker hits with the shower maximum detector hits on a $\Delta\phi = 0.1$ scale. The shower maximum detector is used to assign the track an η coordinate in 0.2 bins. The p_t from the tracker is scaled to match against the calorimeter E_t .

There are two of these crates, one for each half of the detector ($0 < |\eta| < 3.0$). Each crate contains 16 Track/Shower Max Hit Cards (TSMHC), each handling $3.0\Delta\eta \times 0.4\Delta\phi$. The TSMHC cards each receive one central tracker fiber optic cable, one from the intermediate tracker, two each from the barrel and endcap shower maximum detector. The barrel tracker fibers carry four 2-bit p_t 's for four $\Delta\phi = 0.1$ bins. The intermediate tracker fibers carry 1-bit p_t 's. The shower maximum detector fibers have 14 barrel and 16 endcap hit flags covering a $0.1\Delta\phi \times 0.2\Delta\eta$ region. The TSMHC matches the shower maximum detector and track hits on a scale of $\Delta\phi = 0.1$. These hits are assigned an η value from the shower maximum detector and a p_t value from the track. This information is transmitted to the Electron Matching Crates.

Each TSMHC covers eight $\Delta\phi = 0.1$ tracker bins across a region of $3.0\Delta\eta \times 0.4\Delta\phi$. There are 4 $\Delta\phi = 0.1$ bins for each end of the central tracker, and 4 $\Delta\phi = 0.1$ bins for each intermediate tracker end. The contents of these bins are transmitted to the Muon Trigger Correlated Muon Logic Cards (CMLC) as 3-bit hits (2-bit p_t , 1-bit empty/full).

The TSMHC's also transmit to the Electron Match Region Match Cards (RMC) a 3-bit code for each 0.2×0.2 region covered by the TSMHC. The codes reflect the number of shower maximum detector matched tracker hits within these regions and the p_t thresholds of those hits. These are used by the RMC's to correlate electron and tracker hits.

Electron Match Crates

These crates match the shower maximum detector with outer tracker hits and the calorimeter electron candidates. The results are counted up in categories of calorimeter and shower maximum detector matched track agreement on p_t , disagreement on p_t , and shower maximum detector match only with the calorimeter. For each category, the total number of matches above each of 6 E_t thresholds is accumulated. Failures to match shower maximum detector and calorimeter are also counted.

There are eight of these crates: four each for the barrel and intermediate regions of the detector. Each barrel crate covers up to $1.6\Delta\eta \times 3.2\Delta\phi$. It contains 8 Region Match Cards (RMC), each handling $0.8\Delta\eta \times 0.8\Delta\phi$ that use the data from the PRA and TSHMC cards to match an electron in the calorimeter with a shower maximum detector hit and a track passing the appropriate p_t threshold. Each crate has 3 Hit Summary Cards (HSC). The HSC is the building block for hit sums. At the lowest level, the HSC's receive data from 4 RMC's. At the higher levels additional HSC's receive and combine data from 4 HSC's at lower levels.

Energy Sum/Jet Threshold

This crate continues the global E_t , E_x , and E_y sums and adds up the number of 0.8×0.8 calorimeter towers above a series of 8 thresholds. The crate covers both the barrel and intermediate regions of the detector. It contains four Energy Sum Cards (ESC), each handling $1.6\Delta\eta \times 6.4\Delta\phi$. They receive the E_t , E_x , and E_y sums from eight Energy Cards each of which spans $1.6\Delta\eta \times 0.8\Delta\phi$, in four Calorimeter Summation Crates. The output consists of E_t , E_x , and E_y sums which are transmitted to the Total Energy Summation Module in the Final Decision Crate.

This crate also contains four Jet Threshold Sum Cards (JTSC), each handling $1.6\Delta\eta \times 6.4\Delta\phi$, which receive the numbers of 0.4×0.4 towers over thresholds from eight Energy Cards, each covering a region $1.6\Delta\eta \times 0.8\Delta\phi$, in four calorimeter summation crates, and sums these numbers into a series of numbers of 0.8×0.8 towers in two overlapping arrays over thresholds. These results, covering $1.6\Delta\eta \times 6.4\Delta\phi$, are transmitted to the total Jet Threshold Card in the Final Decision Crate.

Muon Trigger Crates

These crates count up the number of muon tracks passing one or two p_t thresholds that are uncorrelated with other trigger subsystem information, are correlated with a track over one of two p_t thresholds in the barrel tracker or over one p_t threshold in the intermediate tracker, and are correlated with a "quiet" calorimeter region.

There are four Muon Trigger Crates. Each crate covers $2.5\Delta\eta \times 3.2\Delta\phi$. Each crate contains eight Correlated Muon Logic Cards (CMLC), each of which covers $2.5\Delta\eta \times 0.4\Delta\phi$. Each CMLC receives one fiber apiece from the barrel, intermediate and forward muon systems. Each fiber carries information for up to 2 tracks from a $1.0\Delta\eta \times 0.4\Delta\phi$ region in the barrel and forward systems and a $0.4\Delta\eta \times 0.4\Delta\phi$ region in the intermediate systems. Per track, there are two p_t bits, one bit for the $\Delta\phi = 0.2$ bin and up to three bits for the $\Delta\eta = 0.2$ bin. Each CMLC also receives a 3-bit hit code (two bits of p_t , one bit empty/full) for the highest p_t track in the corresponding central and intermediate tracker $\Delta\phi = 0.1$ bins from the TSMHC cards. In addition, each CMLC also receives "quiet flags" from the RMC cards, one for each 0.2×0.2 region covered by the CMLC. The CMLC correlates these data into four histograms, each with 2 p_t bins of three bits each.

Each Crate also contains 4 HSC's which sum the histogram outputs from the 8 CMLC's. Each HSC covers one histogram. The output of the HSC's is transmitted to an additional group of four more muon HSC's in the Final Decision Crate.

Final Decision Crate

The Final Decision Crate sums up and makes cuts on calorimeter total and missing E_t . It sums up the total number of $0.8\Delta\eta \times 0.8\Delta\phi$ calorimeter regions over thresholds and makes cuts on this information. It sums up the total number of electron and photon candidates within the categories and p_t thresholds described above and makes cuts on these. It also sums up the total number of muons over the p_t thresholds and makes cuts on these. The resulting trigger decision is then forwarded to the clock and control logic.

There is one Final Decision Crate. It contains three electron HSC's, which accept information from the barrel and two intermediate crate electron HSC's. The information from these electron HSC's include a 3-bit count of the number of calorimeter - track - shower maximum detector matches above each of 6 thresholds where there was agreement between calorimeter and track energy, plus the number of matches without p_t agreement placed above the threshold of the calorimeter value, plus the number of calorimeter - shower maximum detector matches without a track match, and finally a bit that indicates a failure to match calorimeter and shower maximum detector information. This information is sent to the Final Decision Module (FDM) described below.

The crate also contains one Energy Sum Card (ESC) that sums the E_t , E_x , and E_y sums, which are transmitted from one of the Energy Sum/Jet Threshold Crates. The output of the ESC is also sent to the FDM. The crate also contains one Total Jet Threshold Card which reads in numbers of overlapping 0.8×0.8 towers over thresholds from the Energy Sum/Jet Threshold Crates.

The crate also contains four muon HSC's. These complete the summation of the muon hit histograms on the data output from the muon HSC's in the four muon trigger crates. The information in these histograms includes a 3-bit count of the number of muon to tracker to quiet-region matches at each of two p_t thresholds, the number of muon to quiet-region only matches, the number of muon to tracker only matches, and the number of muon only matches. The histogram results are sent to the FDM.

The crate contains one Final Decision Module (FDM) card. It has four separate memory lookup (MLU) tables on it, one each for total and missing transverse energy triggers, jet triggers, electron/photon triggers and muon triggers. Each of these MLU's produces a code that is sent to the Final Decision Logic Card (FDL). The energy sum MLU computes the total E_t , E_x , and E_y , uses the latter two to compute missing E_t , and then compares the total calorimeter global E_t and missing E_t sums against programmable thresholds to produce a 3-bit trigger vector that is sent to the FDL. The jet trigger MLU compares the number of towers over threshold with a series of programmable trigger requirements. The results from the electron match HSC and muon trigger MLU are compared against a series of thresholds.

The FDL card receives the output from the electron, muon, jet and energy MLU's. The FDL contains another series of MLU's that make the final Level 1 trigger decision based on programmable combinations of this data. The FDL card decision is transmitted to the the Global Clock/Control Crate which turns this into a Level 1 Accept if the detector systems are able to read out the event.

Level 1 Clock & Control

Overview

The Global Level 1 Trigger System (GL1) receives a subset of the data from subsystems participating in the Level 1 trigger. The data for each 16 ns bunch crossing, sent by each subsystem every 16 ns, arrives at the input of the GL1 after 36-131 crossings have elapsed since the event occurred, where the fastest time is that of the calorimeter subsystem and the slowest that of the muon system [12]. The GL1 issues a Level 1 trigger Accept/Reject at its output 183 bunch crossings after the event occurred. The functionality and timing of the clock and control system are described in detail in Refs. 12, 13, and 24.

Clock & control operation

The GL1 not only makes the final Level 1 decision, but it also controls the start of any Level 1 readout, disables this readout, and distributes scheduling signals for tests and calibrations. Therefore, all subsystems, even those not involved in the Level 1 decision, must communicate with the GL1.

All Level 1 trigger and DAQ pipelines in the subsystems and the GL1 must be driven with a common clock synchronized to the SSC crossing frequency. This clock is distributed centrally through the GL1. It is phase locked to the SSC machine clock. The processing of the GL1 decision is driven by this 16 ns cycle clock.

Subsystems send trigger data as soon as they are available; they do not perform global event time synchronization. The GL1 must be time align the trigger data it receives since the arrival time for the data of a particular crossing differs among the different subsystems.

The particular bunch crossing number of a subsystem's data is indicated by a local bunch counter, named Bunch Phase. The Bunch Phase is used to check the alignment of the trigger data from each subsystem. The subsystem also counts Level 1 triggers. This number, Level 1 ID, is used to check that the data read out correspond to the correct Level 1 Accept.

Level 1 accept or reject

After the GL1 processing, a decision is made to issue a Level 1 Accept or Reject for each bunch crossing. That signal is accompanied by additional GL1 data. These data are not necessarily propagated to all components of a subsystem, but are used for diagnostic purposes by the central processing of a subsystem. The data include data valid, the Bunch Phase, the Level 1 ID, and several bits of "readout type." The Level 1 ID counts Level 1 Accepts. The bits of "readout type" allow information to be sent to subsystems that may react differently to different trigger types.

Ambiguous and overlapping triggers

If a subsystem cannot identify data with a particular crossing, it sends the same data applicable to a number of crossings with each crossing. Since most triggers rely on a coincidence of components, the result will still be to select a single crossing for every Level 1 Accept. Subsystems determine the number of bunch crossings to store upon each Level 1 Accept. The present design does not allow back-to-back Level 1 Accepts.

Trigger control

Subsystems send trigger control information to the GL1. This includes a "busy" bit to indicate that the subsystem DAQ cannot read out the event, a "fatal error" bit to indicate that the subsystem trigger has had a fatal error and a "trigger data error" bit to indicate that the subsystem trigger data may be corrupt. An example of a trigger busy would be a reset process that had not been completed. Allowance is made for subsystems that cannot complete a Level 1 Reset sequence in a single 16 ns interval. The trigger busy bit prevents a Level 1 Accept being issued to a subsystem that must use several crossings for a reset. An example of a fatal error would be a disagreement between the Level 1 trigger number counted by the subsystem and the number sent by the GL1. An example of a trigger data error would be either a parity check failure or another type of internal data checking failure. The control information also includes the subsystem's crossing number associated with this data. All subsystems active during a run must send this control information to the GL1, whether or not they send data used in the Level 1 decision.

Tests

The operation and scheduling of tests during a run are controlled by the GL1. A schedule of all tests to be run is downloaded to the GL1 before the beginning of a run by the online computer system. The schedule can be modified during a run in response to changing detector status. The GL1 sends to all subsystems a Test Enable signal, with the type of test indicated in the Test Trigger Type bits, a fixed number of crossings before the crossing when the test is to be done. This enables subsystems to prepare for the test. The tests are arranged in such a way as to allow many subsystems to do tests simultaneously. The subsystems set up their tests so that the test data will be contained in the exact crossing indicated by the GL1 Test Enable Signal. The GL1 inhibits normal triggers a fixed time before the test trigger crossing so that there will be no interference. The GL1 then sends out a Level 1 Accept for the appropriate crossing to read in the test data. This accept is accompanied by a bit in the Readout Type set to indicate that it is a test trigger.

Level 1 clock & control distribution

Each subsystem crate in the detector is linked to the Level 1 Clock and Control system through three basic connections: the Clock, Control, and Status cables. The Clock cable distributes the 16 ns crossing clock to the subsystem crate, the Control cable distributes Level 1 trigger information, and the Status cable returns Level 1 status information from the subsystem crate to the Clock and Control system. The Control and Status cables are required to carry a 16-bit Control or Status word for each 16 ns interval, thus having a data rate of 1 Gbit/s.

Figure 8-15 shows the Level 1 Clock and Control distribution system. This system consists of two basic crate types, the Global Clock and Control, and the Level 1 Distribution crates, which together distribute the crossing clock and control information to, and receive status from, all trigger and DAQ subsystem crates.

The Status cables carry information from the subsystem crates to the Global Clock and Control crate. Each Level 1 Distribution crate consolidates the status of as many as 64 subsystem crates, and transmits that on a Status cable to the Global Clock and Control crate.

The Global Clock and Control crate uses the output from the Final Decision crate, along with machine clock and beam pickup information, plus operator control through the DAQ processor system to generate the root Clock and Control cable output. These signals are sent to as many as 16 nearby Level 1 Distribution crates, which in turn rebroadcast the Clock and Control cable signals to as many as 64 subsystem crates each. The signal distribution between and within the clock and control crates is done by Level 1 Clock Control Boards (L1CCB) and Level 1 Clock Control Fanouts (L1CCF). The L1CCB is the basic interface card for all subsystem crates to the Level 1 trigger. The L1CCF is used in conjunction with the L1CCB through a backplane interface to provide a fanout rebroadcast of the Clock and Control cables to crates throughout the detector, and to consolidate returned Status cable data from these same crates.

The function of the L1CCB is to receive the machine clock and Level 1 trigger information on Clock and Control cables via the front panel and rebroadcast this information to L1CCF or subsystem electronics cards across four identical backplane busses. Simultaneously, the L1CCB receives and consolidates subsystem status information from these cards and transmitting it upwards on the Status cable connection on the front panel. For test and development purposes, the L1CCB also can operate in a stand-alone mode in which it emulates the machine clock and Control cable flow from an on-board oscillator and pattern memory.

The L1CCF card is used in the Global Clock and Control and Level 1 Distribution crates to perform the 1-to-4 front panel fanout of Clock and Control cable information received from the L1CCB via the backplane, and the 4-to-1 fan-in of Status cable information received from the front panel connections. It provides the link between the L1CCB in the current crate and the L1CCBs in the next level of crates below.

The Global Clock and Control and Level 1 Distribution crates reside physically near each other in the counting house. They are connected via Clock, Control and Status cables driven electrically. The subsystem

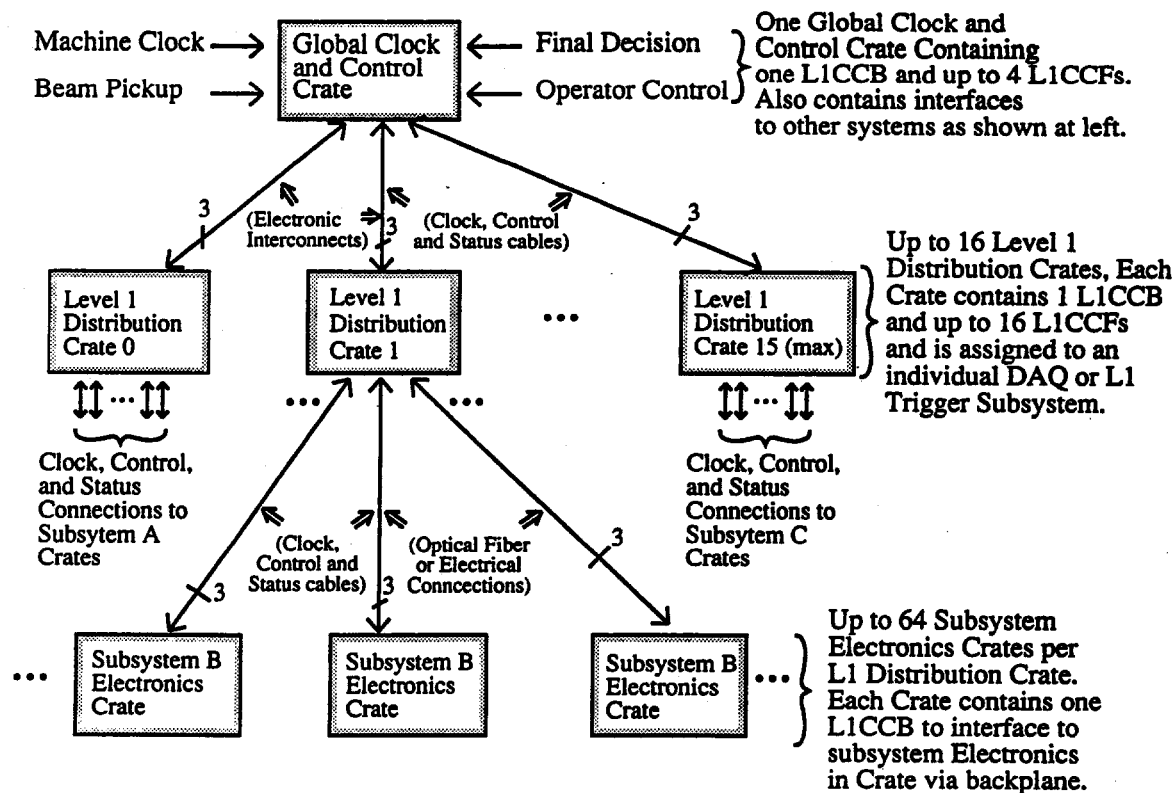


FIG. 8-15. Level 1 Clock and Control System.

crates reside both on and off the detector. Those on the detector, such as front-end DAQ subsystem crates, are connected to the Level 1 Distribution crates using fiber optic Clock, Control and Status cables. Those off the detector, such as Level 1 Decision crates, are driven by electrical cables.

Level 1 pipeline length

The pipeline length is the sum of the times to collect the data from the interaction, generate subsystem information to be used in the trigger calculation (trigger primitives), propagate this information to the global Level 1 trigger decision logic, calculate the Level 1 decision, and propagate this decision back to the subsystem. The critical timing path is determined by the subsystem that takes the longest to send its trigger data to the Level 1 decision logic and the subsystem that has the shortest readout pipeline buffer.

The trigger primitives are formed by the front-end electronics crates. The primitives are subsystem dependent in both their generation time and format. For the calorimeter, shower maximum detector and tracker, trigger primitives are delivered relatively quickly after the interaction. However, due to the long drift times ($\sim 1 \mu s$) in the drift tubes, the muon primitives arrive a comparatively long time after the interaction. This drift time is one of the major contributors to the overall Level 1 Pipeline length.

Another major contributor to the overall pipeline length is the time to transmit, via fiber optic cables, the primitives from the front-end crates on the detector to the Level 1 Decision logic off the detector. With total cable lengths in the 90 m to 115 m range, the resulting transmission time is approximately $0.5 \mu s$. Since this distance must also be traversed on the return path with the Level 1 Accept, the total impact on the pipeline due to the length of the fiber optic cable connection is $1 \mu s$. Location of the Level 1 trigger electronics in the detector hall, where it would be inaccessible during colliding beam operation, would reduce the roundtrip propagation delay by 400 ns.

The Level 1 Decision logic is the third component of the overall pipeline length. The exact pipeline length from receipt of primitives to Level 1 Accept differs in the different front-end subsystems and the type of trigger being formed. Overall, the computation time for the Level 1 Decision from the received

trigger primitives is approximately $1\ \mu\text{s}$.

The final component of the Level 1 Pipeline length is the propagation of the Level 1 Accept/Reject decision back to the front-end crates. This time is the same for all subsystems, as they all use the same Clock and Control circuitry. The contribution of this component is approximately $0.66\ \mu\text{s}$. The total Level 1 pipeline length has been analyzed for each detector subsystem [12]. The muon system requires the longest pipeline length—255 crossings, including a contingency of 25 crossings.

In all cases, the Level 1 Decision is generated 183 cycles after the interaction. For the total and missing energy, jet, and electron triggers, which are associated only with the calorimeter and tracker, this necessitates 72 additional cycles of delay in the front of the FDM card [18]. This delay offsets the long primitive generation time present in the muon system.

8.2.5. Level 2 Trigger function

Introduction

The Level 2 trigger uses flexible, programmable algorithms to confirm and refine the Level 1 trigger by combining data from different subsystems, including new information which was not available at Level 1. For example, Level 2 rejects conversions by matching the silicon tracking information with the central tracker and the calorimeter energy cluster. The full shower maximum detector information is used to improve the electron isolation cut criteria. The muon p_t threshold is sharpened by matching the tracking detectors with muon segments using the full granularity of the outer tracker and muon systems, as well as the information from additional muon layers and the silicon tracker.

The Level 2 trigger selects events based on signatures of high p_t inclusive particles such as leptons and jets, using transverse and longitudinal shower profiles, track reconstruction and segment matching. In addition, topological and geometrical correlations can identify fundamental signatures such as di-leptons and multi-jets.

The Level 2 rejection requirement is determined by the Level 1 trigger rate (10–100 kHz) and the Level 3 bandwidth (1 kHz). A factor of 10 to 100 rejection is required. With a maximum trigger rate of 100 kHz, the average time between Level 2 decisions must not exceed $10\ \mu\text{s}$. For some events, longer decision times are possible (up to $50\ \mu\text{s}$).

The Level 2 is an integral part of both the trigger selection and the data readout stream. Much of the data required by Level 2 will flow from Level 1 or along paths shared with Level 3. Fast signals for synchronization, control, and inhibit are managed in common with Level 1 and the data acquisition. The results of Level 1 processing can be used by Level 2 to isolate “regions of interest” for further processing and to steer the local extraction of information. Level 2 can also select which processing algorithms to apply to an event candidate based upon the results of Level 1. In the same manner, Level 2 can provide a flag of physical topology in order to speed the selection of the primary algorithms for the Level 3 processing farm.

Level 2 tasks

The purpose of the Level 2 trigger is to retain the interesting events selected by the Level 1 trigger, while rejecting most background. To this end, Level 2 executes algorithms which are unavailable to Level 1 because of insufficient input data or insufficient execution time. Furthermore, the Level 2 algorithms is sufficiently flexible to reject unexpected backgrounds that pass the Level 1 trigger.

The Level 2 trigger must reduce the Level 1 trigger rate by a factor of 10–100. Cuts are made using data from tracking systems, the shower maximum detector, the muon system, and the calorimeter. The following are the primary tasks of Level 2:

- Muons: Refining muon selection by matching ϕ and p_t between the muon and tracking systems, by improving muon p_t resolution based on the tracking systems, and using calorimeter information, including, if necessary, isolation cuts.

- **Electrons:** Refining electron selection by matching ϕ between the shower maximum detector and tracking systems, by matching momentum between the calorimeter and tracking systems (E/p), by using isolation cuts, if necessary, and by identifying photon conversions using the silicon tracker.
- **Photons:** Refining photon selection by using the shower profile in the shower maximum detector and by using calorimeter information (HAC/EM).
- **Jets:** Identifying jets using clustering and fixed-cone algorithms.
- **Taus and high- p_t hadrons:** Identifying isolated high- p_t hadrons using the tracking information, the shower maximum detector, and the calorimeter information.
- **Secondary vertices:** Identifying displaced secondary vertices using the silicon tracking system, possibly in conjunction with other tags such as lepton identification).
- **Missing E_t :** Identifying missing E_t and refining missing E_t measurements by summing over calorimeter clusters and by correcting for muon momenta.
- **Lastly, and most importantly,** selecting combinations of multiple electrons, muons, jets, stiff tracks, neutrinos, and displaced vertices in order that a flexible, powerful list of triggers can evolve with the physics needs of SDC.

In addition, other tasks for Level 2, which are justified by past experience, include:

- Rejection of unexpected backgrounds and noise which are not identified by the less flexible Level 1 trigger algorithms.
- Redundancy with Level 1, implemented by replaying the Level 1 trigger selection with more sophisticated algorithms.
- Fast online checking and diagnostics for computing trigger efficiencies.

Level 2 architecture

Our studies of the architecture of the Level 2 trigger have thus far focused mainly upon defining the nature of the Level 2 selection criteria and the input data required to implement these criteria. In the discussion of the Level 2 trigger which follows, we present the data requirements along with a conceptual design of a Level 2 implementation. We anticipate that this design will undergo significant evolution and development as our studies progress. The conceptual design attempts to incorporate well-defined paths for the required data and to define the principal Level 2 processing blocks in a fashion that can exploit rapidly advancing processor technology at the time of implementation. For this reason, the Level 2 trigger will be implemented in a modular architecture with clear, well-defined interfaces between the "modules". The modular structure will allow us to choose the appropriate "tools" for each function at a later time, without changing the overall Level 2 architecture. Each of the modules corresponds to one of four logical steps in the Level 2 architecture:

- Local Level 2 Processors
- Data Collector Buffers
- Global Level 2 Processors
- Level 2 Decision Logic

The Local Level 2 Processors process and format Level 2 trigger data from a local region of the detector, usually from a single detector subsystem such as the calorimeter, into trigger primitives (ϕ , p_t , etc.) which will be loaded into Data Collector Buffers. Dedicated digital processing techniques, such as histogrammers (*e.g.*, as described below for the silicon tracker) and associative memories [25], can be used for this step.

A critical element of the local processing step at Level 2 is extracting the input data to the processors. The technique for providing input data clearly depends on the design of the front-end electronics of each subsystem. In many cases the data bandwidth is quite high; however, the Local Level 2 Processors can use information from Level 1 to identify limited "regions of interest" from which to extract data.

The Data Collector Buffers collect, assemble, and store the Level 2 primitives from the Local Level 2 Processors. These buffers are organized into η - ϕ cells. For each η - ϕ cell there is a word of, for instance, 64 bits which is downloaded in parallel to the Global Level 2 Processors. Well-known techniques such as FIFO registers and fast dual-port memories can be used for the buffers.

The Global Level 2 Processors examine the Level 2 primitives and make a list of interesting candidates for the Level 2 Decision Logic. The highly segmented nature of the two dimensional η - ϕ cells suggests the use of commercial "vector" processing techniques with programmable algorithms. This technology is advancing rapidly; present techniques include associative memories [25], neural networks [26], and massively-parallel single-instruction/multiple-data machines [27], as well as other image processors [28].

The list of candidates produced by the Global Level 2 Processors are combined and processed by the Level 2 Decision Logic in order to extract physical quantities, such as total energies and masses, and to form combinations of quantities such as di-leptons or missing E_t plus leptons. This is an appropriate application for "scalar" processors such as Digital Signal Processors (DSP), RISC processors, or Transputers.

Level 2 data flow

A block diagram of the Level 2 data flow is shown in Fig. 8-16. The detector is segmented into η - ϕ cells with two different structures: axial (ϕ) sectors for the tracking, muon, and shower maximum detector systems, and projective towers for the calorimetry. The Level 2 data flow is organized into two parallel streams of pipelined steps, corresponding to the two η - ϕ structures. The tracking data stream goes to the Track Match processors; the calorimetry data stream goes to the Energy Cluster processors. Then the two streams merge into the Level 2 Decision logic.

The tracking and muon data are sent to the Local Level 2 processors, which calculate track parameters and load these into the Track-Match buffers. The shower maximum detector data are digitized in the front-end electronics and transferred to the dual-port DAQ buffers and to the Track-Match buffers. The calorimetry data are taken directly from the Level 1 buffers (Table 8-1) and loaded into the Energy-Sum buffers. The Track-Match buffer is organized into $1024\phi \times 30\eta$ cells; the Energy-Sum buffer has the calorimetry granularity of $64\phi \times 60\eta$.

Table 8-4
Additional Level 2 trigger information beyond that from Level 1.

Subsystem	Total channels	L2 segmentation	L2 data
Silicon Track	$1024 \phi \times 6 \eta$	$0.00626 \phi \times 0.4 \eta$	5 p_t bits 3 vertex bits
Barrel Track	$1024 \phi \times 2$ ends	0.00625ϕ	5 p_t bits
Int. Track	$1024 \phi \times 4 \eta \times 2$ ends	$0.00625 \phi \times 0.25 \eta$	5 p_t bits
Shower Max	$1024 \phi^\dagger \times 15 \eta \times 2$ ends	$0.00625 \phi^\dagger \times 0.2 \eta$	8 bits energy
Barrel Muon	$1024 \phi \times 10 \eta$	$0.00625 \phi \times 0.2 \eta$	5 p_t bits for 2 tracks
Int. & Fwd. Muon	$1024 \phi \times 8 \eta \times 2$ ends	$0.00625 \phi \times 0.2 \eta$	5 p_t bits for 2 tracks

[†] For $|\eta| > 2.2$, the ϕ trigger tower dimensions switch to 0.0125 and then to 0.025.

The information stored in the Level 2 Track Match buffers is described in Table 8-4. The track and muon data for stiff tracks are extrapolated to a common plane near shower maximum in the EM calorimeter. The p_t information is coded on a five-bit logarithmic scale, sufficient to allow threshold cuts

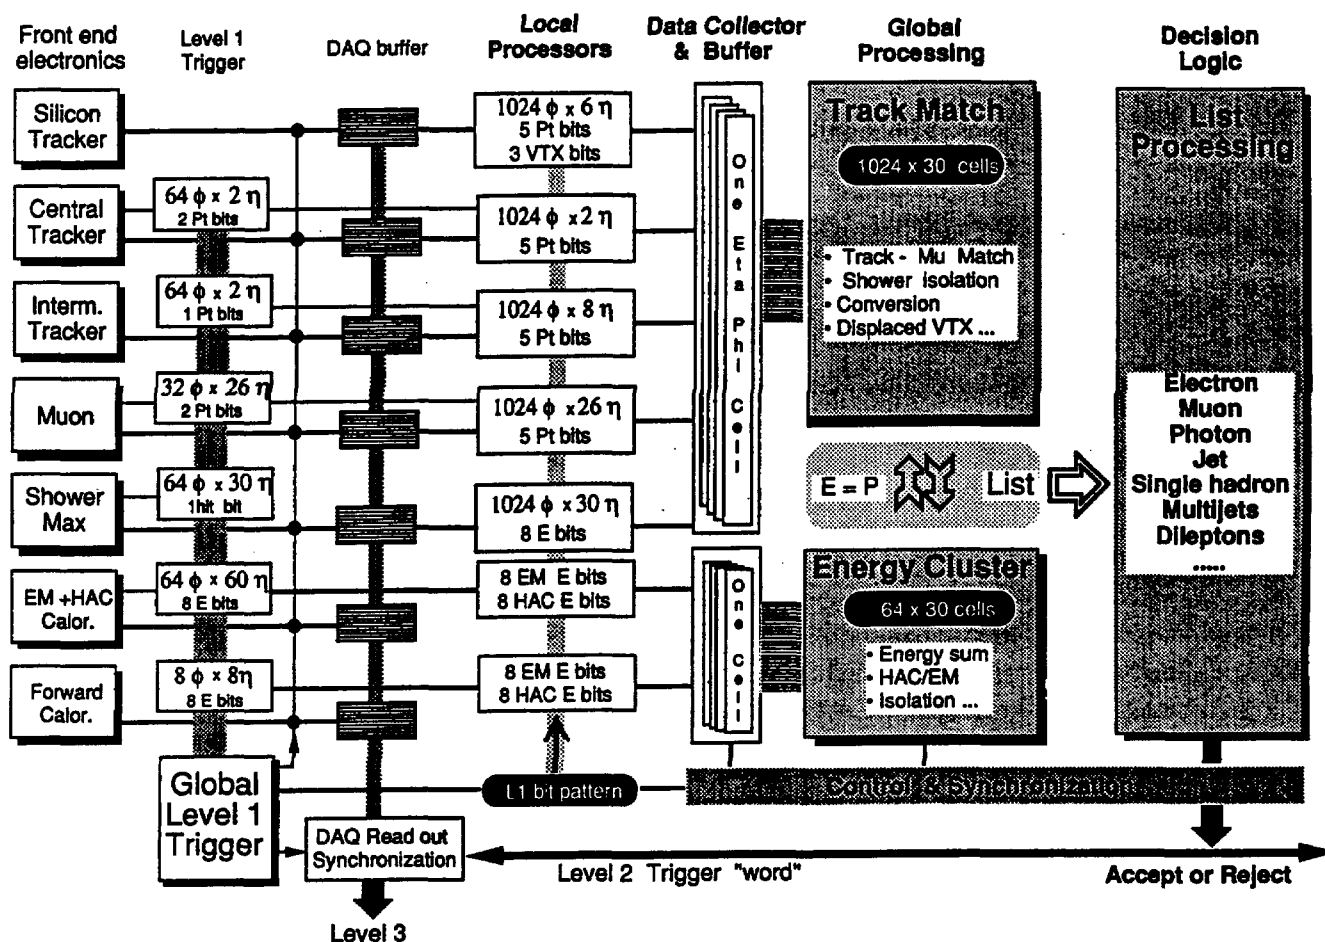


FIG. 8-16. Data Flow in Level 2 trigger.

and p_t matches between 10 and 100 GeV/c with a precision of 20%. In addition, photon conversions and displaced vertices are tagged by vertex bits from the silicon tracker.

From the Data Collector Buffers, data is moved to the Global Level 2 Processors, which do track matching, clustering, and other algorithms in every cell in parallel. Cells that satisfy preselected trigger criteria are added to an output list formed by the processors. The Level 2 Decision Logic operates on this list to complete the Level 2 trigger decision.

After a Level 2 Accept, all the digitized data are sent on to the Level 3 farm by the DAQ system. The entire Level 2 decision process can take up to 50 μ s. Nonetheless, a Level 2 decision comes out of the pipeline every 10 μ s for the maximum input rate.

A Level 2 prototype

A prototype for global Level 2 processing, based on the Associative String Processor (ASP) developed by ASPEX (U.K.), is now under construction. The ASP is a single-instruction/multiple-data machine, in which a single program controller manages the common processing flow of a large number of identical Associative Processing Elements (APE's) arranged in a linear configuration called a "string." A number of VLSI circuits, each containing a number of chained APE's, are linked together to form a complete system. A 25-MHz chip containing 64 APE's, each with a 64-bit data register, has been developed by Hughes Aircraft Company; a 40-MHz chip is under development. Hybrids of 16 ASP chips (1024 processors) will be available soon. A description of the ASP and its possible application to the Level 2 SDC trigger can be

found in Ref. 29.

Calorimetry and track trigger processes involve a mixture of topological (data searching, thresholding, matching, filtering and clustering) and numerical (data scaling, counting, energy accumulation, physics features extraction) sub-tasks. The ASP is highly efficient on the primary topological data reduction phases, where its associative searches, parallel comparisons, and neighbor communications are used at their best. The ASP is also efficient on numerical operations when they can be simultaneously applied to a sufficient number of selected parallel elementary processors. Subsequent low-parallelism Level 2 algorithms, for which the density of active data objects is low and which involve correlations of partial results at distant locations in the string of processors, can be executed more efficiently in a conventional scalar processor in the Level 2 Decision Logic.

Typical Level 2 algorithms have been tested on an ASP simulator, with data mapped onto an array of 128×128 cells, with two APE's per cell, each with a 32-bit data register. In this simulation, all the data from a 0.05×0.05 tower were loaded into the 64-bit register of two neighboring APE's. Two microseconds were needed to find electromagnetic energy clusters or hadronic jets by two simple threshold comparisons. Within $5 \mu\text{s}$, electrons isolated in a single tower could be recognized by comparing the energy deposited in the 8 nearest neighbors with a lower threshold. These preliminary results indicate that the Track-Match and Energy-Cluster algorithms can be performed in less than $10 \mu\text{s}$ with the 40-MHz chips now under development at Hughes Aircraft Company. A compact ASP machine with 8192 APE's on a single 9U VMEbus board is now under construction at Saclay. We intend to test it in the 1993 beam test of the barrel calorimeter pre-production prototype.

The cost for the global aspects of the Level 2 trigger— Data Collector Buffers, Global Level 2 Processors, and Level 2 Decision Logic—was estimated assuming the ASP would be used for the global Level 2 processing. We are confident that the estimated cost will be representative of the cost of the technique eventually chosen.

8.2.6. Level 3 trigger function

Overview

The Level 3 trigger provides the final event selection prior to permanently recording the data on mass storage. This selection is made by applying software trigger algorithms to the events using a large farm of general purpose processors. The use of software algorithms and general purpose processors allows sophisticated algorithms to be employed in making the final event selection. In some cases, these may be the same algorithms used in the offline analysis, a practice that reduces the difficulty in estimating trigger biases. The use of software also provides a tremendous degree of flexibility in tailoring the trigger algorithms to deal with signatures of particular processes, incorporate new trigger algorithms, obtain calibration event samples, and eliminate unanticipated backgrounds.

An event that is accepted by the Level 2 trigger will be read out by the data acquisition system, with the complete event record routed to one of the processors in the farm. This processor will apply the Level 3 selection algorithms to the event to decide whether to keep or discard the event. Since the processors in the farm act in parallel, a large number of events can be simultaneously processed. The Level 3 trigger is designed to accept a Level 2 trigger rate of ~ 1 kHz that may be expanded to 10 kHz. It provides a reduction of ~ 100 in trigger rate. Further information on the Level 3 processors, event delivery to the farm, and Level 3 control software can be found in Section 8.9. In this section we describe how the Level 3 trigger participates in the overall trigger strategy and discuss the various types of trigger algorithms to be used in Level 3.

Level 3 trigger strategy

The Level 1 and Level 2 hardware triggers are designed to make event selections that reduce the event rate to a level that can be reasonably accommodated by the data acquisition system. In addition, the hardware triggers may incorporate trigger elements that would be difficult to emulate in software or would require an inordinate amount of CPU time to implement in Level 3. Once the trigger rate is reduced to a few kHz, it becomes cost-effective to implement in Level 3 those algorithms that require only a few MIPS-seconds of processing time (one MIPS equals one VAX-11/780 unit of processing).

It is essential that the Level 3 processing be restricted, compared to full offline reconstruction, in a way that keeps the average CPU time per event at a low level. We estimate that the average CPU time per event will be 100 MIPS-seconds. That requires the farm to have 10^5 MIPS for a 1 kHz Level 2 rate.

To keep the average CPU time per event low, the Level 3 trigger will first use relatively fast trigger algorithms to quickly reject those events that are easily eliminated. The remaining events will require more sophisticated processing, leading to the application of increasingly sophisticated selection algorithms. Ultimately, a small fraction of the events may be subjected to partial or full offline reconstruction.

Level 3 algorithms

In this section we list some of the algorithms that may be useful to implement in the Level 3 trigger. The purpose of this listing is to give a flavor for the different types of algorithms that may be employed in Level 3. This is by no means a complete or exclusive listing; the flexibility to easily tailor the selection criteria to the needs of the experiment is one of the virtues of the Level 3 trigger.

We first list some fast algorithms that may be appropriate to the initial Level 3 selection. Note that some of these algorithms may be required at an earlier stage of the trigger to reduce the event rate to manageable levels.

- Recalculate various E_T measurements using the full calorimeter resolution to sharpen thresholds
- Use advanced pattern recognition in the calorimeter and shower maximum detector for electron identification
- Make a precise position match between electron track segments and the shower maximum detector profile
- Perform jet clustering and refine jet triggers
- Include calorimeter information in muon selection
- Include forward calorimeter information in E_T sums

The following algorithms are expected to be somewhat slower than those listed above:

- Perform complete 3-D tracking in the muon system
- Match track segments in all tracking layers for stiff tracks
- Incorporate various calorimeter energy corrections, including corrections for dead material, gaps between barrel/endcap calorimeters, and gaps between calorimeter modules

Finally, we list some algorithms that are considerably slower and that will require significantly more CPU time/event:

- Advanced pattern recognition algorithms in tracking system to find photon conversions, secondary vertices, and topology signatures (*e.g.*, τ 's)
- 3-D track finding to locate vertex position, correlate tracks from a single interaction, and reduce ambiguities in tracking
- Track fitting to help solve pattern recognition ambiguities and eliminate fake tracks
- Full offline reconstruction of the event

8.2.7. Schedule

The schedule for the trigger system is driven by the following needs:

- The complete trigger system must be installed by January 1999 to allow sufficient time for completion of system integration and testing before data taking in October 1999.
- The serial integration of subsystems into the trigger system must begin with partial subsets of these subsystems in 1997.
- Operation of the trigger system with the DAQ system must be started in 1996 in order to have trigger and DAQ systems ready for operation with front-end systems in 1997.
- Test stand-alone trigger systems must be available for test-beam running and bench test of front-end systems in 1993.

These considerations lead to the schedule shown in Table 8-5. This schedule affords approximately three years for continued design studies and detailed system design.

Table 8-5
Trigger Schedule Milestones.

Major Milestone	Scheduled Date
Start Final L1 Design	January 1993
Test System Available	May 1993
Complete L1 Design Specs.	December 1993
Final L1 Design Review	January 1995
Complete Design of Level 1	June 1995
Start Final L2 Design	January 1994
Complete L2 Design Specs.	December 1994
Final L2 Design Review	July 1995
Complete Design of Level 2	December 1995
Deliver L1 & L2 Prototypes	June 1996
Initial Delivery of Trigger Interfaces	June 1996
Delivery of Trigger System Begins	January 1997
Begin Integration & Test w/partial systems	June 1997
Begin Integration & Test w/final systems	January 1998
Commission Trigger System	October 1999

8.3. Calorimetry front-end electronics

8.3.1. Introduction/requirements

A key design goal of the readout electronics for the calorimeter is to provide a uniform, cost-effective solution for the barrel, endcap, and forward calorimeter towers and for the shower maximum detectors. Although the various calorimeter channels are similar in that they all need to measure pulse heights at sampling rates equal to the bunch crossing frequency, they differ in the details of their photodetectors, in ionization amplitudes, in their dynamic range requirements, and in the nature and amount of trigger data they must provide to the Level 1 and Level 2 trigger processors.

For the calorimeter tower channels in the barrel and endcap regions, the input signals are provided by photomultiplier tubes (PMT's). To optimize the dynamic range, the maximum linear pulsed current from the PMT should be matched to the maximum energy deposition in the detector. A reasonably good PMT can deliver ~ 1 nC in a $100 \text{ ma} \times 10 \text{ ns}$ pulse. When combined with the number of photoelectrons/GeV expected from the calorimeter, these numbers fix the PMT gain at a few $\times 10^3$. This indicates the use of a PMT with 6 to 7 stages. The PMT signals are very fast, with peaking times less than about 10 ns, half-width less than about 24 ns, and duration less than 100 ns. A signal from a shower in a prototype electromagnetic tower is illustrated in Fig. 8-17. The time response is consistent with expectations for the fluors in the wavelength shifter, and it indicates the approximate time structure which we expect. The details of the input signal shape will depend upon the fluors chosen for the scintillator and wavelength shifter. The size of the input signal is at least 250 photoelectrons per incident GeV. Fast integration or pulse shaping is required to preserve the advantages of the scintillator speed and to minimize the effects of pileup from crossings near in time.




FIG. 8-17. A typical signal from a shower in a prototype electromagnetic module, with and without signal clipping.




FIG. 8-18. Physical layout of calorimeter front-end electronics.

The dynamic range of the readout system must be at least 2×10^5 (18 bits) to meet the calorimeter requirements for a minimum detectable signal of 20 MeV E_t (> 5 photoelectrons at 90 degrees) and a full-scale signal of 4.0 TeV E_t . The minimum 20 MeV signal is specified so that 100 MeV side leakage from electron showers can be detected. Another benchmark for the minimum detectable signal is the energy deposition of a minimum ionizing particle which is ~ 250 MeV in the electromagnetic section and 1.4 GeV in each hadronic section. The 4 TeV E_t cutoff is chosen since a jet that deposits this amount of energy in one of the calorimeter towers occurs rarely.

The phototube signals must be sampled at the beam crossing frequency of 60 MHz, and there must be the capability to read out multiple samples per event or to combine (with weights) samples nearby in time. The system must continuously accept and process analog signals from the detector and simultaneously

output digital event data in a deadtimeless fashion. It must also have good accuracy, linearity, and crosstalk properties. While satisfying these stringent requirements, the electronics must have low power consumption and must not degrade the overall performance of the calorimeter.

In addition, calorimeter towers must be grouped into trigger towers of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$. These tower signals must be digitized at the beam crossing frequency with a dynamic range of at least 12 bits and then transmitted to the Level 1 trigger processor. A nonlinear 8-bit digitization is sufficient.

The electronics system must provide capability for three methods of calibration: by radioactive source, by electronic charge injection, and by optical pulser. The calorimeter design includes a remote-controlled radioactive source which visits each tile in the calorimeter. The PMT anode currents resulting from the source must be integrated, digitized, and monitored. Ability to inject a calibration pulse with programmable amplitude and time of arrival must be provided. This charge injection system must be able to exercise the full dynamic range of the readout. It will be used to check for missing codes, establish scale overlaps, perform cabling checks and other diagnostics, as well as to calibrate the electronic response. The source and charge injection systems provide each wedge of calorimetry with a self-contained calibration system. Electronic calibration should be accurate to 0.5%, as should the monitoring of the anode currents from the calibration sources. The final linearity check of the optical and readout systems will be performed using a flasher system whose optical attenuation is controllable over many orders of magnitude.

Each $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ region of the barrel calorimeter is initially instrumented with six electronics channels, consisting of four transverse electromagnetic towers of 0.05×0.05 in one longitudinal segment and of two longitudinal hadronic segmentations (HAC1 and HAC2) of 0.1×0.1 . Each 0.1×0.1 region can be upgraded to 10 readout channels by regrouping of waveguide fibers and addition of more photomultipliers until there are two longitudinal electromagnetic compartments (EM1 and EM2), each with four transverse towers of 0.05×0.05 . The endcap is configured with the two electromagnetic compartments on day one. In the barrel and end-cap calorimeters combined, there are $\sim 13,500$ electromagnetic calorimeter channels and ~ 6300 hadronic compartment channels before upgrade. After the full upgrade in segmentation, there will be $\sim 20,700$ electromagnetic channels and ~ 6300 hadronic channels. For triggering, these channels are summed into ~ 3200 towers, each with separate EM and HAC readout. The number of trigger towers remains the same when the calorimeter is resegmented. The segmentation near $\eta = 3$ in the endcap is somewhat different. The readout of the forward calorimeter will be adapted from the readout design for the central calorimetry.

The physical layout of the front-end electronics for both the calorimeter towers and shower maximum detectors is shown in Fig. 8-18. Signal processing electronics is mounted in the photomultiplier base or on a nearby circuit card. Front-end boards packaged in crates mounted on the calorimeter provide signal buffering, and in some designs, digitization. Trigger tower energy sums are also formed on the front-end boards. Trigger and readout data are transmitted to the trigger and data acquisition electronics in the counting house at ground level by fiber optic links. The front-end electronics will be subdivided into stand-alone subsystems, each supporting a calorimeter module of a size convenient for independent beam testing during commissioning. Besides their primary function of data acquisition, these subsystems will provide convenient means of checkout, calibration and defect location. In addition, the subsystems will be able to trap event data on-the-fly during normal data acquisition, store these data, and make them available for diagnostics. Each subsystem crate receives the set of fast timing signals from the trigger system via a Level 1 Clock Control Board which then distributes the signals to each front-end board.

Current architectural approaches

In order to address the design challenge of readout with large dynamic range at high sampling frequency, we are currently pursuing two different architectural approaches. They both address the same design requirements, but satisfy them in significantly different ways. One approach, described in Section 8.3.2, is called the Switched Capacitor Array (SCA) system and provides a 2×10^5 dynamic range. The other, called the Digital Phototube Readout system, is described in Section 8.3.3 and proposes a 20-bit dynamic range to achieve a least count of 10 MeV and a full-scale response of 10 TeV. Both designs meet the calorimeter requirements.

In both approaches the photomultiplier signal is sampled at the bunch crossing frequency with the appropriate digital event data made available to the DAQ system after a Level 2 Accept trigger. In both approaches, the phototube information from all crossings is stored without prior zero-suppression. (In other words, the storage is not data driven as it is in some silicon strip and wire chamber readout schemes.)

Both architectural approaches can deliver the required data to the trigger processor, and both must interface to the data acquisition and clock and control systems in the same way. They also share the same basic physical layout, as shown in Fig. 8-18. Some electronics are mounted in the photomultiplier base or on a nearby circuit card; front-end boards contain data storage and digitization circuits; the front-end boards are packaged in crates mounted on the calorimeter wedges; trigger tower formation and digitization is performed on the front-end boards; and trigger and readout data are transmitted by fiber optic links to the trigger and data acquisition electronics in the counting house at ground level.

The two readout approaches differ primarily in whether they digitize before or after storing the signals in Level 1 and Level 2 buffers. The SCA approach is based on processing the phototube signals through a preamplifier and shaper, storage of analog samples in a SCA, and digitization subsequent to a Level 2 Accept trigger. The Fast Digitizer approach is based upon a floating-point flash encoder which splits the current of the phototube signals into an array of gated integrators, the outputs of which are then digitized at the beam crossing rate by a flash ADC.

The details and present status of each of these approaches are described below. Each approach offers its attractions and disadvantages. Each is confronted by somewhat different technical challenges. While tests of individual components are encouraging, neither approach has yet fully demonstrated that it can deliver the requisite dynamic range in a system implementation. In light of the importance of the calorimeter system, we are continuing the development of both approaches until one proves adequate.

8.3.2. Switched Capacitor Arrays readout system

Introduction

In the Switched Capacitor Array (SCA) approach to calorimeter readout, the design goals are addressed by a system based on high-speed arrays of switched capacitors, implemented in custom-designed CMOS integrated circuits. Analog-to-digital conversion occurs after the Level 2 trigger decision, where data rates are significantly lower, in order to reduce power requirements of the digitization. The requisite dynamic range is provided by two overlapping signal amplitude ranges.

The design is also influenced by the requirement to implement sensitive low-noise analog circuitry on the same integrated circuit where a significant amount of fast digital switching occurs asynchronously to the analog input signals. The potential for noise pick-up is severe in such situations. Efforts are therefore being made to minimize the number of digital control lines and, where possible, to use balanced digital signals.

Front-end architecture

The SCA system follows the generic front-end architecture scheme described in Section 8.1, with the analog-digital conversion situated late in the signal processing chain, after the Level 2 storage. The basic scheme for a single channel is illustrated in Fig. 8-19. With the ADC in this position, the Level 1 and Level 2 storage pipelines are required to store information in analog form. The amplified and shaped analog signal from a PMT is sampled in synchronism with the beam crossing frequency and stored in elements of an SCA—a different storage capacitor element for each crossing. Once stored, the analog sample remains in the capacitor element until read out or over-written. Lists of address pointers maintain the status of the sample in each capacitor array element. The sample will be read out if it belongs to an event that passes the Level 2 trigger criteria or will be discarded if its event fails either Level 1 or Level 2 criteria. The samples that are read out are then digitized and passed to the Data Acquisition (DAQ) System.

The address list manipulation adds slightly to the complexity of the channel electronics, but results in the very significant advantage that analog information is not transferred from one pipeline to another. The amplitude resolution of the SCA is enhanced by this scheme, since reading and re-storing the analog signals would raise the inherent noise level of the system. A single address list processor services up to 192 channels; the address lists of all channels are thereby manipulated identically.

The basic components of the signal processing chain are:

- Preamplifier and shaper with charge calibration facility.
- Event data processing chain, with dual range splitting ($\times 1$ and $\times 32$), analog Level 1 and Level 2 storage buffers (“pipelines”).
- Analog-digital conversion after Level 2 triggering, and multiplexed digital event data readout.
- Address list processor for managing the SCA analog storage functions.
- Trigger data processing chain for summing and digitizing signals representing towers for Level 1 trigger processing.

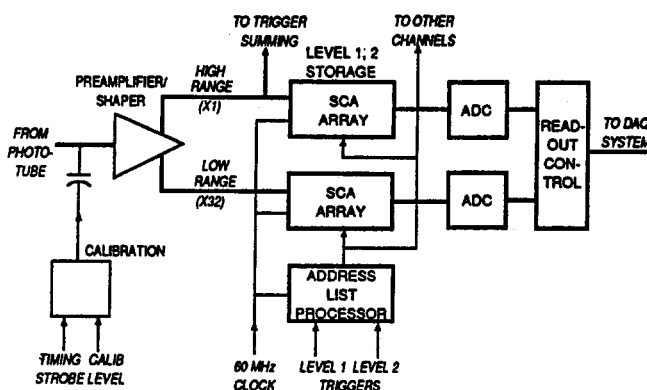


FIG. 8-19. Conceptual architecture of a channel of calorimeter front-end electronics. The PMT signal is preamplified and shaped, split into high and low-range channels, and then stored during the Level 1 and Level 2 trigger in a 256-element SCA analog memory. After a Level 2 Accept, the analog data are digitized and read out.

Block diagram and description

The front-end system is composed of 96 identical front-end blocks, 32 on each end of the barrel and 16 on each end-cap calorimeter. A front-end block, shown in Fig. 8-20, services either one barrel wedge or two endcap wedges. There are two main parts. The first is a group of PMTs and preamplifiers located on Preamp/HV/Calib cards that are mounted on the outside of the calorimeter. The signals from these cards are transmitted over cables to the second part, which is housed in a standard DAQ readout crate, located in the closest available area within the SDC structure. The crate accommodates a number of SCA cards, a Level 1 Clock Control Board, a DAQ CPU card, and a Data Collector and Buffer Memory card.

Each front-end block processes signals from approximately 1000 channels of the calorimeter, and develops two streams of digitized data; one for the trigger processor, and one for the DAQ system. During commissioning or maintenance periods, individual blocks can be operated as autonomous subsystems, independently of other blocks or other front-end systems. The system design is described in detail in Ref. 30.

Light signals from the scintillators are detected by the PMTs, preamplified and delay-line shaped. The shaped analog signals are transmitted as high-range/low-range pairs of analog signals to an SCA card over shielded twisted-pair cables. On the SCA card, the analog signals are sampled at 60 MHz and stored in the SCA ICs. The identity of each sample is maintained in the address controller in terms of the beam crossing number from which the datum originated.

The electronics on the SCA card responds to a Level 2 Accept by selecting the correct SCA sample, digitizing it and transferring it to an output buffer. Data from all the output buffers on all SCA cards in the crate are transferred to the Data Collector and Buffer Memory card and queued for transfer to the DAQ system.

The SCA card also provides trigger information to the trigger processor on each crossing.

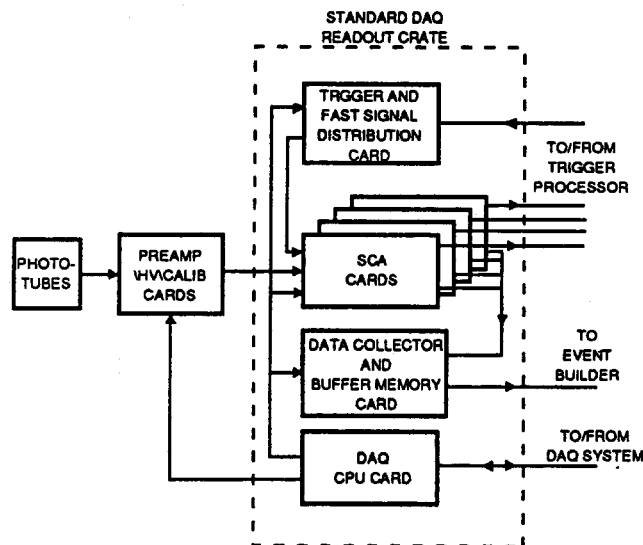


FIG. 8-20. A front-end "block" consists of one crate and a group of Preamp/HV/Calib cards. Signals from a group of cards are transported over cables to the readout crate. The crate contains the electronics that store the analog signals, generate and transmit trigger data, and process and transmit the event data to the DAQ.

Preamp/HV/Calib cards

Groups of ten PMTs are physically supported on Preamp/HV/Calib cards, which also contain preamplifiers, signal shapers, calibration facilities and circuits for supplying the PMT operating voltages. Each card is physically mounted close to the scintillators that it services.

The fast signal from the anode of each PMT is coupled to the input of a fast preamplifier. Because of the wide dynamic range of the signals, each preamplifier has two analog output ranges. These overlapping ranges are separated by a factor of 32 in gain. The high-range signal covers the entire amplitude range, 0 to 4.0 TeV. The low range will saturate for signals greater than about 200 GeV. The preamplifiers have a 16 ns delay-line shaping network common to both ranges. For calibration, a charge signal of controllable amplitude is selectively injected into the input of each preamplifier. The controllability of the amplitude of injected charge will be adequate to measure the linearity of each channel to better than 0.5%.

A separate current integrating amplifier is also connected to each PMT anode. It is used to monitor tile radiation damage by measuring the average anode current while a radioactive source of known strength is exposed to the scintillator. The integrated currents are processed by a separate ADC system that is readout by the Control/Monitoring Network. The results of this measurement will be used to establish the operating voltage of the PMTs. This measurement technique is expected to be accurate to 0.5%.

The Preamp/HV/Calib cards also provide for individually controlling the gain of each of the ten PMTs. The gain of each tube is controlled by varying the high voltage applied to it. A common high voltage is distributed on the the Preamp/HV/Calib card to the high-voltage circuit for each tube. These circuits consist two elements connected in series—a voltage divider (that supplies voltages to each dynode of a PMT) and a voltage regulator. Most of the high voltage appears across the voltage divider, the remainder across the voltage regulator. The amount of voltage absorbed by the regulator (and hence the amount applied to the divider and the tube) is controlled by a separate DAC for each tube. Under command of the Control/Monitoring network (Section 8.9), each DAC generates a control voltage that is applied to its associated regulator.

SCA cards

A functional block diagram of the SCA card is shown in Fig. 8-21. The prime function of the SCA card is to receive, process and digitize the signals from 192 analog channels, i.e., 96 PMTs. As many as 18 SCA cards can be mounted in an SDC standard DAQ crate. Initially, there are 3 cards in a barrel crate and 3 cards in an endcap crate. The shower maximum detector electronics shares the crates.

Timing and strobes

The SCA card receives the set of fast timing signals from its crate's Trigger and Fast Signal Distribution card via backplane wiring. The set of timing signals includes: Level 1 Accept, Level 1 Reject, Level 2 Accept, Level 2 Reject, the 60 MHz System Clock and Reset. The timing of these signals is stable to ± 0.5 ns with respect to the beam crossing. Because the propagation delays from the source of these signals to the particular SCA card will differ from card-to-card, a means of correcting the propagation delay is provided on each card. These delays are remotely controllable via the Control/Monitoring Network.

Input buffer

The input buffer receives the differential analog signals from the Preamp/HV/Calib cards. Calibration pulses of controllable amplitude can be inserted into selectable channels at this point also. The amplitude of the calibration pulses and the selection of channels to receive them is remotely controlled via the Control/Monitoring Network.

The input buffer furnishes copies of the high-range analog signals to the trigger summing function. The summing preserves the timing and shape of the input signals to allow association of the trigger information with a particular beam crossing. The later stages of the trigger processing are thereby able to receive a sample of the tower activity each and every beam crossing.




FIG. 8-21. A functional block diagram of the SCA card. The prime function of the SCA card is to receive, process and digitize the signals from 192 analog signal channels.

SCA and ALP circuits

Each SCA card contains 12 SCA integrated circuits (IC's). Each SCA IC has 16 input and 16 output channels; the SCA card therefore accommodates 192 channels. (In this context, a "channel" services either a high-range or a low-range signal.) The SCA IC can sample and store samples of the analog signal at its input at a 60 MHz rate.

The SCA IC is an externally addressed (steered) analog memory. It contains 4096 sample and hold analog memory cells subdivided into 16 channels of 256 cells each. Each cell stores the sample of an analog signal from one beam crossing. The SCA IC is a dual-port device; it can simultaneously store and retrieve analog samples from two different cells (see Fig. 8-22). By means of addresses externally generated by the Address List Processor (ALP) function, each cell can be assigned to one of four states: Idle, Level 1 Pending, Level 2 Pending, and Output Pending. Internally, the ALP maintains four lists of cell addresses, one list for each of the four cell states.

In response to a reset from the trigger, the addresses of all cells are assigned to the Idle list; the other three lists are empty. After reset, as each analog sample for each beam crossing is stored by the SCA, the address of the cell holding that sample is placed in the Level 1 pending list. A Level 1 Accept causes the address of the cell holding the sample from the appropriate beam crossing to be placed in the Level 2 Pending list. A Level 2 Accept causes the appropriate cell address from that list to be placed in the Output Pending list. The contents of cells whose addresses are in the Output Pending list are sequentially transferred to the ADC for digitization. Addresses of cells rejected by Level 1 or Level 2 Rejects or whose contents have been digitized are returned to the Idle list.

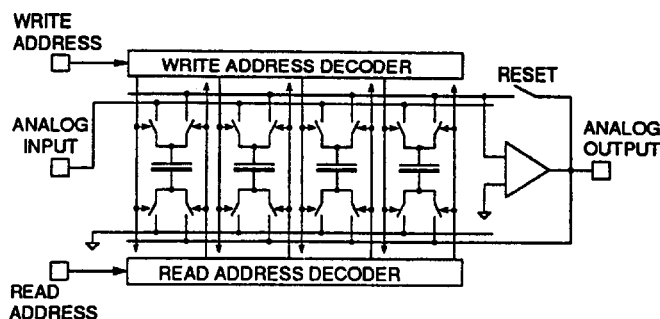


FIG. 8-22. A partial diagram of the dual-ported analog memory cell structure of the SCA integrated circuit.

Analog-to-digital converter

Each SCA card contains 12 ADC integrated circuits for performing analog-to-digital conversion of event data. Each IC contains 16 12-bit analog-to-digital converters and 16 digital FIFO buffers that can store the results of the last four analog-to-digital conversions. All 192 ADC channels on the card perform their conversion simultaneously. A conversion takes $32 \mu\text{s}$. The digitized values are transferred under control of the Readout Control function to the Data Collector and Buffer Memory card. This transfer takes place over a bus on the crate's backplane.

Trigger summing and digitization

Pre-processed data from the calorimeter are required by the trigger processor to generate Level 1 and Level 2 triggers. These data consist of the sums of analog signals representing the energies deposited in individual elements of towers of 0.1×0.1 . The sums are generated in the Trigger Summing function on the SCA card. The span of the analog input signal is equivalent to a transverse energy range of 300 MeV to approximately 1 TeV in the calorimeter. Twenty nonlinear 8-bit ADCs digitize the analog output of each of the trigger sum junctions at a 60 Mhz rate. Each SCA card transmits 20 digital data streams, representing 20 towers, to the trigger processor. Present plans are to implement 2 streams per 1 Gbit/s optical fiber link.

Trigger and Fast Signal Distribution card

The Trigger and Fast Signal Distribution card receives the set of fast signals (listed above) from the trigger system and distributes them via the crate's backplane to all other cards in the crate.

Data Collector and Buffer Memory card

The function of the Data Collector and Buffer Memory card is to gather the event data from all SCA Cards in a crate via a high-speed backplane bus and to transmit it to the Event Builder on an optical fiber. The optical fiber interface and its data stream conform to the SDC DAQ standards.

The card supports other important functions including: trapping an event for analysis by the local CPU, storing but not forwarding data for later use, and sending the Event Builder a synthesized event for diagnostic purposes.

Local CPU card

The design of the Local CPU card is supplied by the DAQ system, and is described in Section 8.9. In the calorimeter crate, the CPU will be able to execute a number of tasks. Among these are:

- Provide the control to operate the front-end block;
- Perform calibration of each channel of the front-end block;
- Perform diagnostics to locate defects;
- Monitor selected event data as it passes through the Data Collector and Buffer Memory;
- Down-load constants to remotely controlled functions;

Status of design efforts

The development of IC building blocks for the SCA-based calorimeter front-end electronics is being carried out using CMOS integrated circuit technology to achieve a high level of integration [1,2,3,31]. Work carried out so far is described below.

Preamplifier/Shaper

An integrated CMOS high dynamic range preamplifier/shaper for scintillation calorimetry has recently been fabricated. A bipolar version based on an AT&T technology has been designed and simulated.

Switched Capacitor Array IC

The most critical component on the SCA card is the Switched Capacitor Array (SCA) integrated circuit [32,33,34]. Implementation of the SCA has required extrapolation of IC designs from those heretofore available. To date, several successful iterations of design, implementation and testing have been carried out. For example, an SCA chip with full simultaneous read/write capability has been fabricated at ORBIT Semiconductor in a 1.2 micron double-metal double-poly technology. The dimensions of the chip are 4 mm \times 6 mm.

The present SCA IC design has 16 channels. Each channel has 256 analog storage locations, and can therefore store up to 4 μ s of analog information at the system clock rate of 60 MHz. The number of storage locations per channel can be easily extended beyond 256 by re-design.

A diagram showing an example of the basic SCA architecture is in Fig. 8-22. Each analog input is internally connected to a bus that distributes the input signal to all 256 storage locations in the channel. Each storage location (sample-and-hold cell) consists of a complementary CMOS transmission gate (complementary switch) and a 0.7 pF double-polysilicon capacitor. The capacitors use a high quality silicon-oxide dielectric of 70 nm thickness. An externally supplied reference voltage is applied to the bottom plate of all capacitors. The voltage stored on each sample and hold capacitor then corresponds to the difference between the input signal and the applied reference at the time its sample and hold switches are opened. During the readout sequence, each capacitor is placed in the feedback path of the reconstruction op-amp. There is one reconstruction op-amp per channel.

In order to achieve 12-bit dynamic range: a) the pedestal variation of each of the storage elements in a channel must be less than 1 mV over its 4 V full-scale voltage range; and b) the time constant for charging each capacitor in the array must be just under 2 ns. It is very desirable that the variations in gain and offset between the various storage locations be sufficiently small that they may be ignored. If this is achieved, it becomes unnecessary to determine calibration constants for every cell in the array. In bench tests of the version implemented in 1.2 micron CMOS, having 256 cells per channel and a 1.5 ns charging time constant, the cell-to-cell variation at a 60 MHz acquisition rate was 0.6 mV rms. This includes variations from both systematic pedestal offsets and from incoherent sources. The useable dynamic range in this version was determined to be from zero to 5 V. These achieved parameters provide a dynamic range of 8000:1 or 13 bits. The circuit can also be used in a differential mode which effectively doubles this dynamic range to 14 bits.

The non-linearity of the SCA circuit is dominated by the voltage-dependent parasitic capacitance of the readout switches. This non-linearity is measured to be ± 4 mV or $\pm 0.01\%$ of full scale, and $\pm 0.5\%$ of value. The measured parameters of the SCA are summarized in Table 8-6.

Table 8-6

Measured performance parameters of a simultaneous read-write switched capacitor array integrated circuit. The array has been implemented in 1.2 micron CMOS, has 16 channels, and 256 storage cells per channel. Measurements were performed at a 60 MHz acquisition rate.

Analog Store Parameters	Achieved Performance
No. channels/chip	16
No. of elements/channel	256
Power consumption/channel	10 mW
Non-linearity	$\pm 0.1\%$ (full scale) $\pm 0.5\%$ (value)
Cell to cell variation	0.6 mV rms
Charging time constant	1.5 ns
Channel dynamic range	8000
Maximum sample rate	90 MHz
Maximum readout rate	200 kHz
Input voltage range	0-5 V
Capacitor droop rate	0.1 mV/ms
Output settling time (0.02%)	1.2 μ s

Address List Processor

A sophisticated Address List Processor (ALP) IC for managing the storage and retrieval of analog samples in the SCA IC has been fabricated and is currently being tested [35]. It contains 4 FIFOs, based on dynamic memory cells, and the arbitration logic required to update the four lists of cell states (Idle, Level 1 Pending, Level 2 Pending, and Output Pending) each 16 ns. This is a straight-forward digital design implemented in CMOS.

Analog-to-digital converter

A linear analog-to-digital converter (ADC) with 12-bit resolution has been fabricated in a 2 micron double-poly CMOS technology [36]. This ADC uses a simple single-slope conversion technique and therefore can be implemented in a small silicon area and exhibits low power consumption. The chip contains a voltage-ramp generator to provide a reference voltage that is a linear function of time and a comparator to detect the ramp-crossing time. An internal counter counts pulses from an internal clock. The contents of the counter are transferred to a latch at the ramp-crossing time. The counter and ramp are shared by all ADC channels implemented in the chip. The analog-to-digital conversion requires 4096 clock cycles. Conversions of 12-bit resolution can be accomplished at a 40 kHz rate. The usable analog input range is within 200 mV of either power supply rail.

This integrated circuit successfully combines high speed CMOS digital frequencies in excess of 60 MHz, with analog signals smaller than 1 mV. Crosstalk and substrate coupling effects have been successfully minimized in this mixed analog and digital chip. A 16-channel version of this chip implemented in 1.2 μ m CMOS is presently being tested.

Photomultiplier tube base

A prototype PMT base containing the series-regulated high-voltage divider has been constructed and tested. The preamplifier is presently being integrated into the tube base and will be tested shortly.

VMEbus implementation of SCA card

In parallel with the IC development, we have begun to study system issues by constructing the first in a series of VMEbus-based circuit cards that will emulate the functions of the SCA Card. Particular concerns to be investigated are the acquisition of low-level analog signals by the SCA in the presence of the 60 MHz driven write-control signals, the random request/reject of previously stored data by the Level 1 and Level 2 trigger systems, and the on-board/backplane noise generated by the transfer of the data to the DAQ.

On this prototype card, the Address List Processor is partially simulated in random logic. It performs simultaneous reads and writes. The parallel ADC system is emulated with analog multiplexers and a single fast ADC. The SCA used on the prototype card is the 16-channel, 256 storage location per channel versions implemented in 1.2 micron CMOS. The prototype board has successfully captured phototube pulses and has been read out over VMEbus. The pedestal variation was measured to be 1 mV, corresponding to a resolution of 12-bits.

8.3.3. Digital Phototube Readout system

In the Digital Phototube Readout system, photomultiplier signals are digitized immediately at the base of each PMT at the beam crossing frequency. The digital event data are then transmitted to the rest of the system on individual flat cables.

The main strength of this system is the fact that all analog operations are localized to within a few inches of the PMT. The analog noise and crosstalk problems which are the traditional bane of detector electronics are thereby minimized. The system is expected to provide a dynamic range of 20 bits with 8-9 bit accuracy. The system design also benefits from the conceptual simplicity, the flexibility and the upgrade capabilities offered by rapidly advancing digital technology. The details of the design exploit the PMT as a nearly ideal current source to accomplish the requisite dynamic range by splitting the PMT current into a number of binary weighted outputs, each of which can be processed with reduced dynamic range requirements.

Overview

As shown in Fig. 8-23, the system has two main components. The first is a PMT Base/Digitizer module which resides on the base of each PMT. It provides the high voltage to the PMT and contains a custom BiCMOS IC which integrates and digitizes the PMT anode current at a 60 MHz rate.

The second main component is a Digital Readout/Trigger card. Groups of these cards are housed in crates mounted within the detector volume. For the current SDC calorimeter design, each card receives signal cables from 18 calorimeter channels spanning three trigger towers of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$, i.e., 12 EM channels, 3 HAC1 channels, and 3 HAC2 channels. It performs the necessary operations to develop data for the trigger processor, to perform the Level 1 and Level 2 storage, to output event data to the Data Acquisition (DAQ) system, to perform calibration of each channel, and to control the high voltage for each PMT.

A single flat cable between each PMT Base/Digitizer assembly and its Digital Readout/Trigger card carries all digital event data to the card, and power and control signals to the assembly.




FIG. 8-23. Block diagram of the Digital PMT readout system.

Calibration

Three methods of calibration are provided—by radioactive source, by electronic charge injection and by optical pulser.

The SDC scintillator calorimeter design [37] includes a remote-controlled radioactive source which visits each tile in the calorimeter [38]. The PMT anode currents resulting from the source are integrated by op-amps in the PMT Base/Digitizer assemblies, digitized by multiplexed ADCs in the Digital Readout/Trigger cards, and monitored by the CPU in the readout crate. This provides each wedge of calorimetry with a completely self-contained calibration system.

Each PMT Base/Digitizer assembly has the ability to inject a calibration pulse with an individually controlled amplitude and time of arrival. This charge injection system must be able to exercise the full dynamic range of the digitizer. It will be used to check for missing codes, establish scale overlaps, perform cabling checks, etc.

The final linearity check of the system will be performed using a flasher system similar to that used by CDF [45], Zeus [16] and many other experiments. The optical attenuation of flashers is controllable over many orders of magnitude.

PMT Base/Digitizer

The PMT Base/Digitizer module mounts in the PMT base. Its circuit diagram is shown in Fig. 8-24. The essential element is a custom Digitizer ASIC. The Base/Digitizer module also includes a Cockcroft-Walton HV supply, high voltage control and monitoring, and calibration circuitry. Test patterns can be generated in the PMT Base/Digitizer module to exercise the complete system at full operating speed. We are targeting a power consumption of 1 W for this module.

Digitizer ASIC

The cornerstone of the Digital PMT Readout system is the circuit which provides high dynamic range digitization of the PMT anode signal every 16 ns. It is designed to produce a floating-point numerical output with 8-9 bit accuracy and 20-bit dynamic range. It is the only analog element in the signal processing chain. A block diagram is shown in Fig. 8-25. This circuit is currently being developed and additional details are given in Ref. 9.

Roughly speaking, the circuit can be described as a hybrid between a conventional “gated integrator” ADC and an “Auto-ranging” Digital Voltmeter. It has six major sections which are described below: a current splitter, a gated integrator/switch, a comparator/latch, an encoder, an analog multiplexer, and a flash ADC.

The PMT input to the Digitizer ASIC is DC-coupled. Resetting the digitizer sampling capacitors between each sample storage avoids problems due to saturation of the input integrator during a burst of large pulses. This results in a DC-coupled system with little or no memory of the energy depositions from previous crossings.

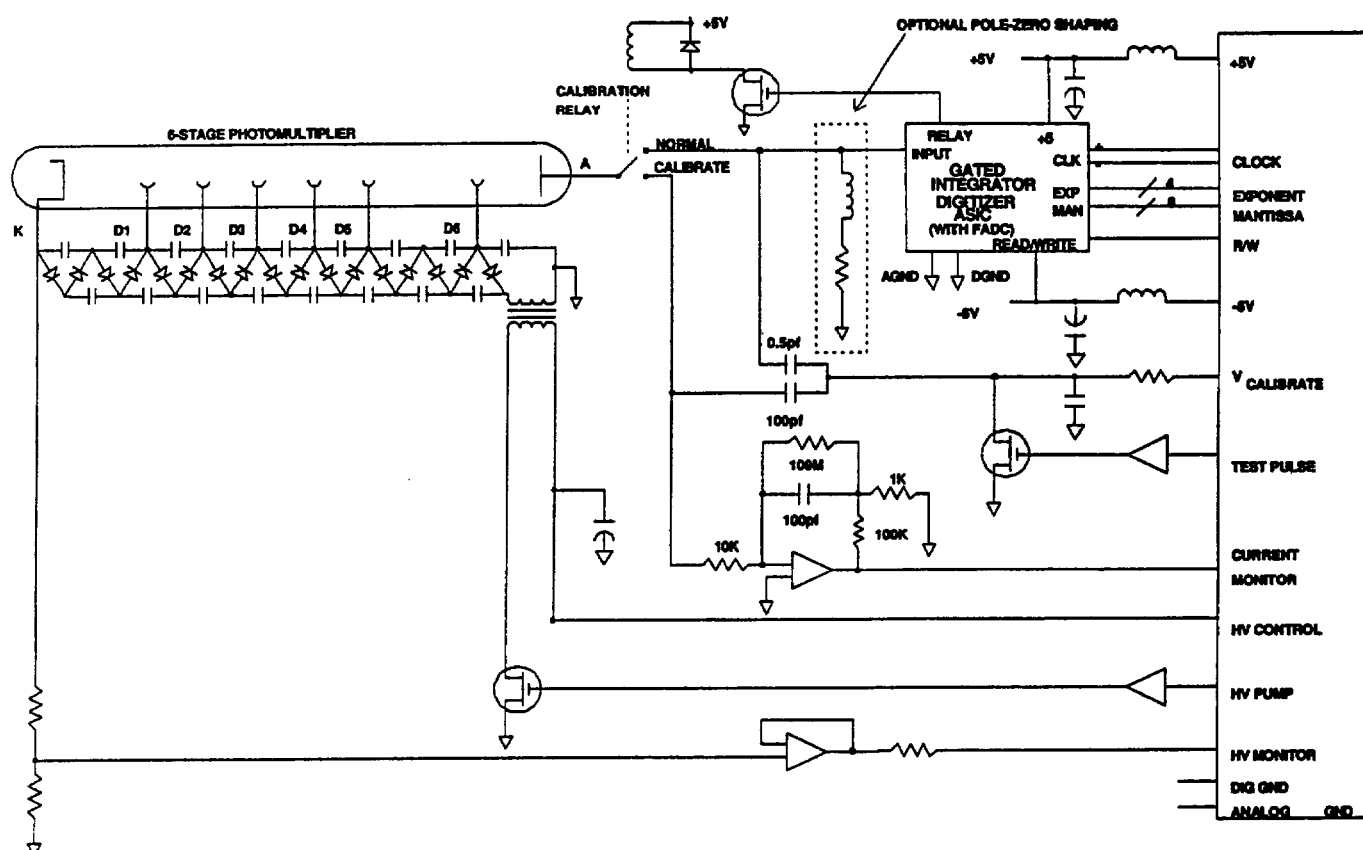


FIG. 8-24. PMT Base/Digitizer Assembly circuit diagram.

Current Splitter

In the Current Splitter section, the PMT anode current is split into approximately ten binary-weighted current outputs. The outputs are sent to identical gated integrator circuits, each of which must handle a modest dynamic range (8-10 bits). These outputs generate the binary scales corresponding to the exponent of the floating point output.

The Current Splitter is the only element in the circuit which has to handle the entire 20-bit dynamic range. It must accurately split the incoming PMT anode current into a number of binary-weighted outputs, while maintaining good (<1%) linearity and fast rise times over the full range of input current (0-100 mA). All other sections of the circuit need to operate with only 8-9 bit accuracy.

Gated Integrator/Switch

In the Gated Integrator/Switch section, each of the approximately ten split PMT currents is sent through a four-way current switch to one of four integrating capacitors of $\sim 1\text{pF}$. Each capacitor has a reset FET and a clamp diode to handle overloads. Each capacitor goes through a four-clock sequence in which it integrates, settles, is sometimes digitized, is buffered to a FADC, and is reset. The sequence proceeds in round-robin manner so that one capacitor is always integrating, i.e., the circuit is deadtimeless.

The integrator is set up for -2V full scale. The 8-bit FADC is set up to digitize voltages between -1V and -2V . Thus, the signal processing accuracy required by the gated integrator is 9 bits. We are designing and testing the gated integrator for 11-bit accuracy in order to take advantage of lower-power 10-bit FADCs becoming available [42].

The range of current which must be passed linearly by an individual gated integrator varies from about 0.1 mA to 0.2 mA . The largest overload pulse which must be handled properly is $\sim 1\text{ mA}$; overload currents larger than this are shunted away at the current splitter. Thus, the dynamic range over which the individual gated integrators have to operate is rather small.

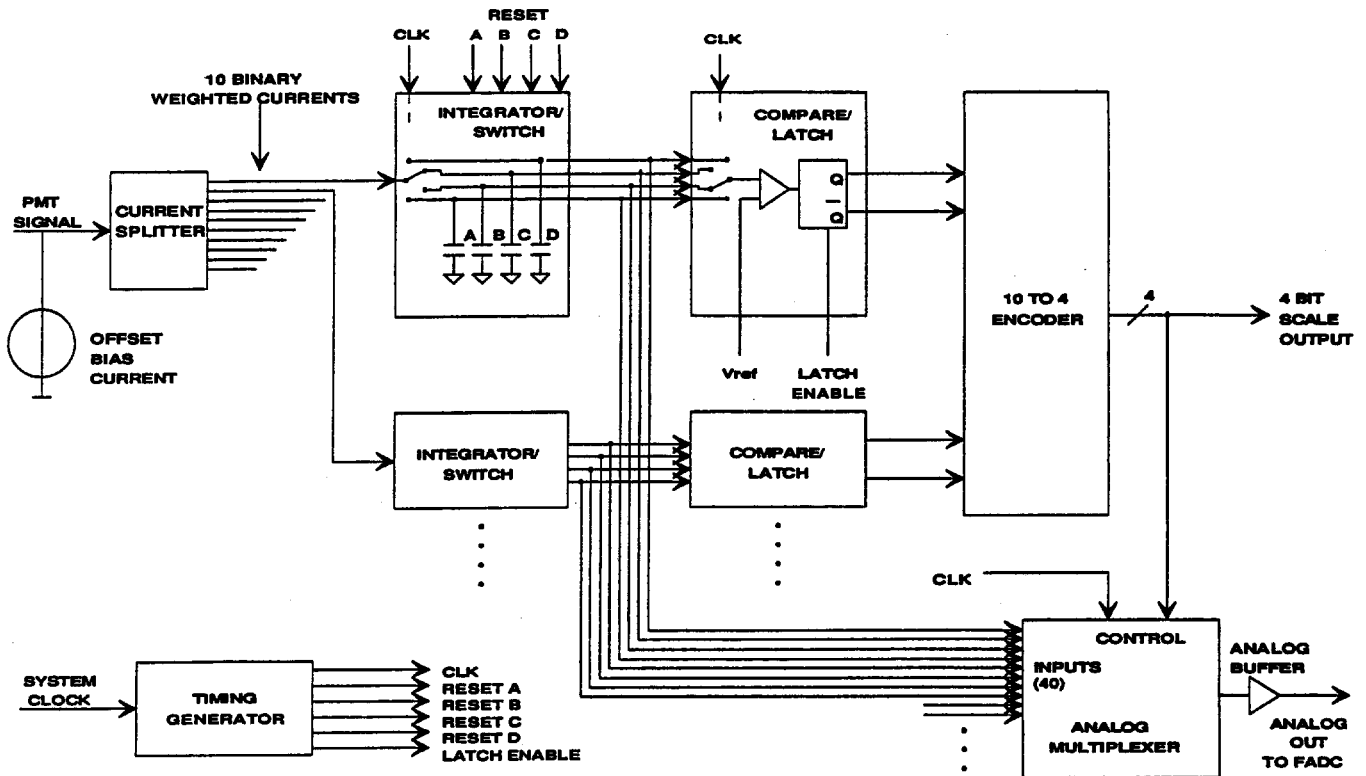


FIG. 8-25. Block diagram of the fast Analog-to-Digital Converter.

Comparator/Latch

The purpose of the Comparator/Latch logic is to determine which of the set of four integrating capacitors has the sample that should be digitized. This determination is made by a set of comparators, each measuring the voltage on a capacitor. The “interesting” sample (having between 1/2 scale and full scale voltage) is on the first capacitor whose comparator has not flipped. This function is analogous to that of an “auto-ranging” digital voltmeter.

Encoder

The Encoder section turns the comparator bits into a 4-bit binary-encoded output. These 4 bits form the binary scale range (exponent) of the floating-point output.

Analog multiplexer/buffer

The Analog Multiplexer/Buffer section is directed by the encoder to select the voltage from the “interesting” capacitor. It buffers that voltage off-chip to the 8-bit Flash ADC (FADC), which produces the 8-bit mantissa of the floating-point output.

Flash ADC

In the initial implementation of the system, the FADCs will be commercially available units. The minimal FADC required is a 60 MHz, 8-bit device, with the reference ladder set up to digitize voltages between -1V and -2V .

FADCs are available from a variety of manufacturers; their performance is continually improving. For example, 10-bit 70 MHz full-flash ADCs are currently available for $\sim\$100$ [42]. Ten-bit devices based on two-stage conversion techniques have clock frequencies approaching 60 MHz and offer much lower power dissipation. Studio cameras for High-Definition TV equipment will require 72 MHz clock frequencies, ensuring a competitive commercial marketplace for these devices.

In the final design it may make sense to integrate the FADC onto the digitizer. FADC sub-circuits are becoming available as standard cells on full-custom ICs. There are also a number of attractive design options which can take advantage of the pipelined nature of our conversion, including half-flash designs, and "one-comparator-per-bit" type designs.

Cockcroft-Walton high-voltage supply

The PMT Base/Digitizer assembly contains a Cockcroft-Walton voltage multiplier HV supply for each PMT. HV wiring is restricted to the metal enclosure of the assembly; the stored energy in the HV system is small. The Cockcroft-Walton HV base [43] has been used successfully in several experiments. If carefully designed and tested, it is a low-power (<100 mW), safe, and reliable means of providing HV to the phototube. The main challenge in using it at SSC energies comes from the low PMT gains required by the high beam energy and luminosity. Low PMT gain requires high-gain electronics, which are sensitive to noise pickup during switching transients from the Cockcroft-Walton. In order to address this issue we have fabricated several test circuits, and have settled on a design (a self-resonant "Class E" converter) which provides sine-wave drive to the C-R bridge and which promises very high efficiency and low noise.

Digital Readout/Trigger card

Digital data are delivered in floating-point format from the PMT Base/Digitizer to the Digital Readout/Trigger Card, which resides in a standard SDC data acquisition crate mounted on a calorimeter module. This card applies calibration constants to the incoming data stream, stores the data in Level 1 and Level 2 buffers, and develops trigger primitives. It delivers calibrated data to the DAQ readout and delivers trigger primitives to the central trigger system. It also provides power and control of the calibration and high voltage functions of the PMT Base/Digitizer, monitors the PMT high voltage, and provides diagnostic capabilities. It also contains circuitry to perform computer-controlled adjustment of the sampling strobe timing on individual phototubes. This permits automated correction for cable lengths, time-of-flight, and propagation delays. Most of the functionality of the card is built into two digital ASIC's: a Digital Pipelined Storage Chip and a Floating-Point Adder ASIC. As shown in Fig. 8-26, the Digital Readout/Trigger card delivers readout and monitor data and receives control information from the "VMEbus-like" backplane of the DAQ crate. A 12-bit DAC provides reference voltages for calibration and HV control to the PMT base. High voltage and current monitor signals from the PMT base are digitized by a 12-bit ADC. Long sequences of Monte Carlo events can be stored in the front-end digital RAMs and "played back" to verify trigger and DAQ performance. It is estimated that the power dissipation will be less than 400 mW per channel.

The Digital Pipelined Storage Chip contains a calibration lookup table, Level 1 "pipelined" storage, and FIFO's for Level 2 and DAQ storage. It also provides the calibrated 12-bit floating-point channel energy to the Floating-Point Adder ASIC. It contains differential cable transceivers and various calibration and test features. It will be fabricated in 1.2 μm CMOS.

The Floating-Point Adder ASIC sums calorimeter channels into trigger towers with $0.1 \times 0.1 \eta - \phi$ granularity. Electromagnetic and hadronic energies are summed separately. Because the sums are digital, they maintain the full precision of the digitization performed in the PMT Base/Digitizer assembly. The trigger tower data are transmitted to the global Level 1 trigger processor on optical fibers.

The Floating-Point Adder ASIC is sufficiently general that it can also be used to form sums of outputs from other Floating-Point Adder ASIC's, *e.g.*, forming total E_t and missing E_t signals. Each ASIC can sum the energies from up to eight inputs. The adders use the same floating point format (8+1 bit mantissa, 4-bit exponent) that is output by the PMT digitizer. They will be fabricated in a BiCMOS process.

In this system, the trigger data are corrected by the calibration lookup table in the Digital Pipelined Storage Chip, using the same set of calibrations as used in Level 3 and offline computing. Thus, the trigger processor is provided with the same digital event data which are transmitted via the DAQ system to archival storage, and the trigger threshold curves will be ideally sharp with respect to the offline algorithms. Since the trigger data are processed in digital form, the performance of the trigger can be well modeled in

Monte Carlo detector simulations and the trigger performance can be verified, on an event-by-event basis, from the energies observed offline.

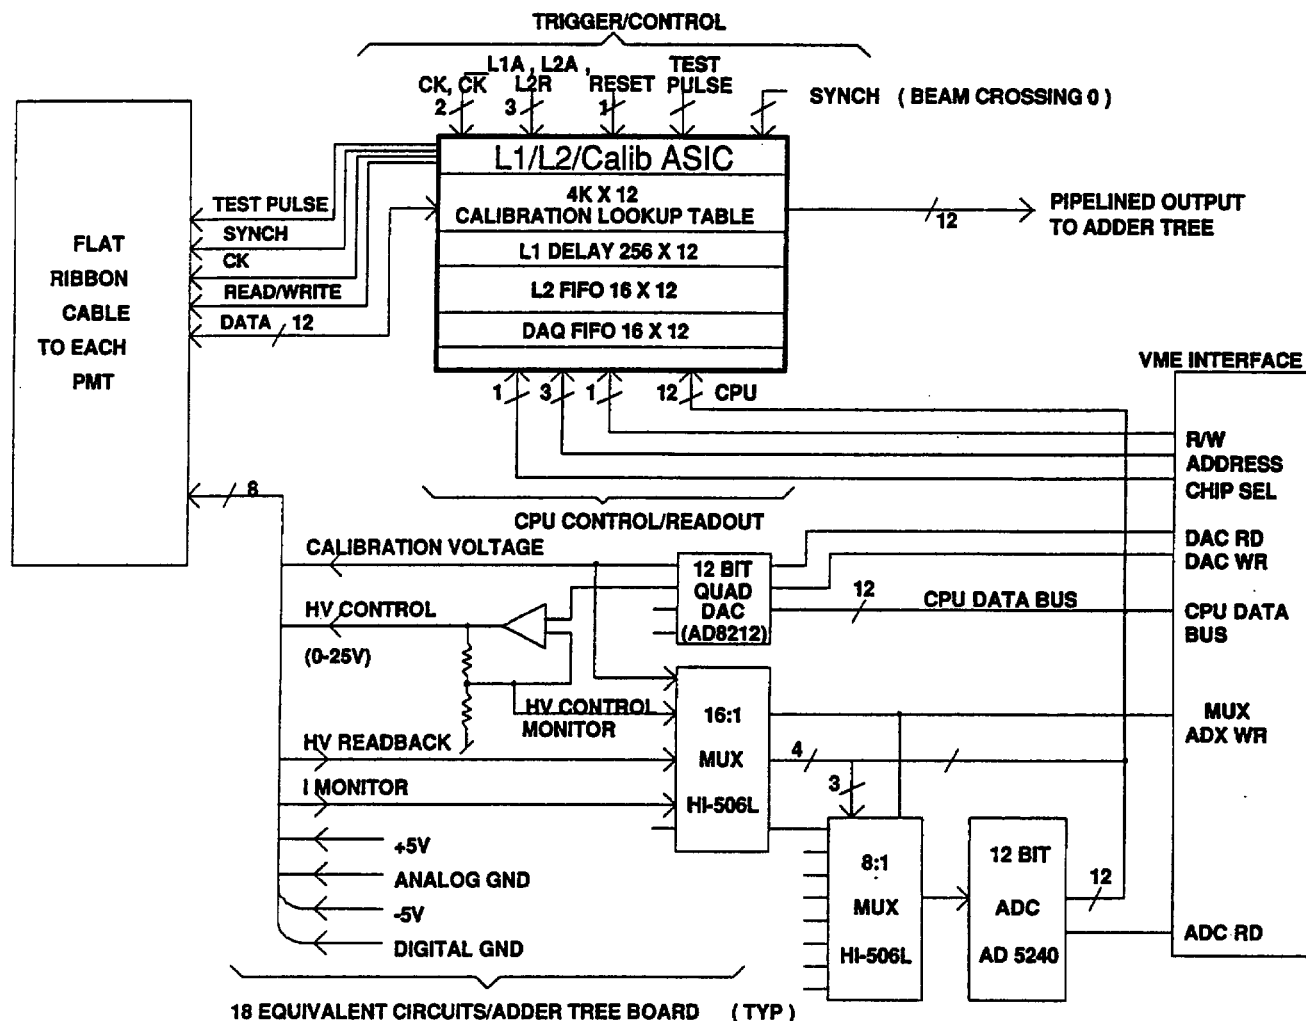


FIG. 8-26. Details of the logic associated with each channel on a Digital Readout/Trigger card.

Status of design

Digitizer ASIC

A first version of a test chip for the current splitter/gated integrator has been fabricated by Orbit Semiconductor. The circuit contains a 9-output splitter, a current switch, integrating capacitor, reset switches of various geometries, and analog multiplexer and output buffer circuits. Measurements of dynamic range, linearity, noise, temperature stability, run-to-run variation in parameters, and output settling time were in good agreement with SPICE predictions. Test results [41] for this initial test IC are summarized in Table 8-7. The splitter was tested with currents over the range 0 to 100 mA. Its accuracy was measured to be $< 0.7\%$ for DC currents. Linearity was confirmed for pulses, and propagation delays less than about 1 ns were measured. The noise performance was measured from the pedestal distribution obtained during repetitive cycling of the acquisition sequence on a single capacitor: capacitor reset, integration, output settling, and measurement. Thus, this measurement includes pedestal sigma from splitter current noise, from timing jitter in the gate width, as well as from digital noise associated with generating the timing sequence. The measured RMS noise was < 0.3 fC ($30 \mu\text{V}$ on a 11 pF capacitor), which corresponds to

0.6 least count when translated to a 20-bit dynamic range. Settling time measurements were made on the rudimentary analog multiplexer and output buffer on the gated integrator/splitter test chip. The tests indicated settling to 8-bit accuracy within 9ns, and 11 bit settling (1 mV) within 13 ns. This exceeds the accuracy required.

Essentially all of the crucial analog issues have been successfully addressed with this chip. A second version of this chip has since been fabricated; a third version has been designed. Each version includes additional features. For example the third version includes on-chip timing and control logic for the gated integrator and output multiplexer, and uses low-level differential signals for digital clocking.

Floating-Point Adder ASIC

Two versions of Floating-Point Adder ICs have been produced in 1.2 μm CMOS. The first version was fully functional in fall-through mode at 60 MHz. The second version contained improvements in the clocking structure, some logic changes due to a redefinition of our floating-point format, and functions in synchronous pipelined mode at ~ 100 MHz.

This design is essentially complete, and seven copies of this sub-circuit will be laid out and fabricated on a single die in order to form the 8:1 adder tree required for trigger tower sums in the Baseline Trigger design.

Digital Pipelined Storage chip

A Digital Pipelined Storage chip was designed and built in 2 μm CMOS. This "Serial-In/Random-Out" (SIRO) chip is representative of the type of digital structure that will be needed for digital Level 1/Level 2 storage. Writing and reading speeds of 56 MHz and 46 MHz, respectively, were achieved.

Future versions of this IC will be fabricated in 1.2 μm CMOS, and will implement the full Level 1/Level 2 storage protocol, differential flat-cable transceivers, and various calibration and test pulse features.

System tests

Over the past year, a number of analog and digital test boards have been built and tested in order to demonstrate the capabilities of the Digital PMT Readout system. In the next few months a full system test will be performed. This will include floating-point digitization based on custom ICs, Cockcroft-Walton HV bases, a digitally pipelined energy sum trigger, and 2-level storage with computer readout. As a first step, several of these have been prototyped using standard components.

As a "proof of principle" demonstration that board-level clocking noise and analog-digital crosstalk will not be a problem for the project, a circuit containing a phototube clip-amp, with a commercial 60 MHz FADC and flat cable line drivers located ~ 3 cm away on the same board has been constructed. The amplifier had a 100 MHz bandwidth and an input sensitivity of 6000 electrons/FADC count. FADC output flat cables were observed digitally and histogrammed. The observed noise level was < 0.5 count, which corresponds to $1/2$ count on the most sensitive scale for a 20-bit dynamic range. On the basis of these tests, it is concluded that crosstalk will not be a problem.

The system tests of the Digital Readout/Trigger logic requires programmable streams of 60 MHz bit fields, and the ability to generate and respond to clock and trigger signals. A board to test these requirements has been built and tested.

Boards have been produced to receive pipelined digital cables and hold the data for CPU readout. Two versions of these boards have been built and debugged, one using custom storage ASICs and the other using commercial parts.

In order to study coherent trigger noise when multiple PMTs are summed, we have produced a pipelined trigger summation card. It accepts output flat cables from 8 PMT digitizers and contains an adder tree which produces a single pipelined energy sum. A second version of this board will be made based on the floating-point adder tree custom ASIC.

Table 8-7
Digital PMT Readout: Gated Integrator/Digitizer Performance Summary

	Specification	Measurement
Least Count (~noise level)	10 MeV (1 fC) = $0.1\mu\text{A} \times 10\text{ ns}$ current pulse from PMT	< 0.6 fC pedestal RMS for gate times < 50 ns on 1st version of Current Splitter/Gated Integrator Test chip
Full Scale	= 10 TeV (1 nC) = 100 mA \times 10 ns pulse = 10^6 least counts (20 bits)	Current Splitter measured linear for pulses up to 100 mA
Full Scale	= 10 TeV (1 nC) = 100 mA \times 10 ns pulse = 10^6 least counts (20 bits)	Current Splitter measured linear for pulses up to 100 mA
Accuracy	1-2% before calibration	Current splitting accuracy measured < 0.7% on multi- ple ranges up to 100 mA
Rise Time	~5-8 ns (faster than scintillator rise time)	< 4 ns over entire dynamic range 0 \rightarrow 100 mA
Time Slewing vs. Pulseheight	<< 16 ns PMT Gate Width	~1 ns difference in splitter propagation delay for pulses of ~0 \rightarrow 100 mA
Gated Integrator		Functional
Capacitor Reset		Functional
Output Multiplexer and Buffer		Functional
Settling Time of Analog Output	< 16 ns to accuracy required by FADC	9 ns settling to ~8 bit accuracy
Temperature Stability	< 0.25% over 10 $^{\circ}\text{C}$ operating temperature range	< 0.5% change in current splitter accuracy for 30 $^{\circ}\text{C}$ change.
Immunity to Run-to-Run Process Variation	< 1% before calibration	Current splitting absolute accuracy measured < 0.7% for test chips from two separate ORBIT runs
System Noise	$\sigma \leq 1$ least count 6,000 e RMS	$\sigma \sim 5000e^-$ (0.8 fC) measured on 60 MHz test board including FADC noise, discrete amplifier noise, digital noise from clock and line drivers, etc.

A major cost driver in the system is the assembled cost of the PMT base/digitizer assembly. To reduce assembly costs a method has been developed for producing the PMT base/digitizer assemblies, literally a hundred at a time, using pick-and-place equipment and a minimum of hand labor. The technique is to build up the base assembly from a series of "disc and washers" made from multilayer PC boards with surface-mount components.

Tests were performed [44] on twelve candidate cables for use between the PMT and Readout crates. The cable selected for use was an 0.025" pitch flat cable with integral ground plane. Signal rise times were preserved (<4 ns) over the ~ 5 m cable lengths required. Differential mode crosstalk was of order 2%. The cable is available in a fire resistant, halogen free jacket at reasonable cost. This is the only cable required in the system.

8.3.4. Shower maximum detector front-end electronics

Introduction

The shower maximum detector (SMD) provides fine segmentation within the EM calorimeter. The 8 ϕ -strips and 8 η -strips per tower are used to measure the position and the shape of the EM showers. Besides formatting read out of data through the DAQ system, data are also generated for use by the Level 1 and Level 2 triggers. This same electronics will also be used to instrument the Massless Gap at the front of the barrel calorimeter.

Requirements and constraints

The SMD electronics process 47,104 SMD electronic channels (94,208 after upgrade) plus the 10,368 Massless Gap electronic channels.

The readout scheme performs the following functions:

- Provide pulse-height information with 12-bit dynamic range and a least count of one photoelectron (half minimum ionizing particle).
- Distinguish between events every 16 ns.
- Supply a single bit per 16 channels to the Level 1 trigger every 16 ns.
- Store analog pulse heights during the 3-4 μ s Level 1 latency.
- Provide information to the Level 2 trigger.
- Read out selected events to the DAQ system at up to 10^4 Hz.
- Provide a read out organization and modularity optimized for the physical granularity, electronics packaging, trigger segmentation, DAQ integration, testability, monitoring, and reliability.
- Survive for 10 years at the radiation levels present behind the calorimeters.

Overview

The functionality of the shower maximum detector front-end electronics is shown schematically in Fig. 8-27. Logically and physically, the electronics consists of four pieces: the photodetectors, the analog signal processors, the front-end digitizers, and the data collection boards.

Shower maximum detector photodetectors

The light signal from each shower maximum detector scintillator strip is detected by an individual photodetector. Several possible pixel photodetectors have been investigated in order to meet our requirements at minimum cost. The photodetectors currently under active study are the multi-channel photomultiplier tube (MCPMT), arrays of avalanche photo-diodes (APD), and a hybrid APD-MCPMT.

The baseline adopted here is the Philips XP1722 64-channel MCPMT. This device is based on an established technology with a gain of 5×10^5 . The gain uniformity is acceptable, but the cross-talk is about 5%. Improved devices with cross-talk below 2% and with green-enhanced photocathodes are being developed with Philips and with Hamamatsu.

The principal alternative is the APD array. APD's have higher quantum efficiency than conventional photocathodes (60% at 520 nm *vs.* 10% for the green-enhanced photocathodes), but the gain is low (100 to 300) and is sensitive to temperature. Therefore, low-noise, high-gain, fast preamplifiers are required, and the temperature must be stabilized. Commercial APD's produced by EG&G (Vaudreuil, Canada) gave encouraging results in the beam-test of a prototype SMD. Prototype APD arrays are being developed for us by Hughes Aircraft and by Advanced Photonix. Long term stability of many channels and engineering aspects are the key issues of such technology.

The APD-MCPMT would use an APD array mounted in a proximity-focussed image tube with a conventional photocathode. The gain is high because the photoelectrons are accelerated in a strong electric field before impinging on the APD. The APD arrays under development at Hughes and at Advanced Photonix would be incorporated in this hybrid device by Litton Electron Devices.

Visible Light Photon Counters (VLPC's) from Rockwell which potentially offer high gain (10^4), high quantum efficiency ($> 80\%$), no cross talk and insensitivity to magnetic fields are currently being developed for scintillating fiber tracking applications. However, their use in the SMD readout requires a better dynamic range (~ 3000 today) and a detailed study of integration into the SMD front-end electronics chain.

Taking into account the uncertainty about the final choice of the photodetector technology, the SMD front-end electronics will be designed to be compatible with any one of our future decisions. However, for the following discussion we assume a 64 channel MCPMT.

Analog signal processors

The output signal from each photodetector channel will be coupled to an analog signal processor consisting of a charge integrator, shaper, gain adjustment amplifier, and discriminator. This device will be packaged into a 16-channel ASIC. The required characteristics are as follows:

- Preamplifier rise time less than 5 ns.
- Peaking time less than 16 ns.
- Recovery time less than 1 μ s for a maximum of 5% occupancy.
- Dynamic range 4000 (compatible with a 12-bit SCA).
- Gain and sensitivity optimized for a MCPMT output signal between 15 fC and 60 pC.
- Gain adjustment of a factor of 5 between MCPMT channels.
- Programmable discriminator threshold with a range of a factor 10 and a logical "OR" of the discriminator outputs for use in the Level 1 trigger.

Additional features, including current integration capability for source calibrations and a variable-amplitude charge test signal, will be also provided. This device is currently under design.

Each front-end analog module will handle 64 channels of analog signal processing from a single MCPMT (32 ϕ -strips and 32 η -strips) and will be connected directly to the MCPMT. This module will contain the MCPMT, the analog signal processing ASIC's, output drivers, and all associated control logic. Slow control and monitoring tasks could be implemented using a low-level microprocessor remotely-controlled through a

low-speed serial bus or point-to-point link. There will be seven front-end analog modules per barrel wedge and five per endcap sector for a total of 768 modules. The 768 front-end modules will be distributed around the calorimeter, allowing relatively easy access to the analog electronics and to the MCPMT's. The analog outputs of each front-end analog module will be connected to a front-end digitizer card. The digital outputs will be connected to the trigger.

Front-end digitizer

The front-end digitizer card will contain signal storage, digitization, and data output to the DAQ and Level 2 trigger systems for 128 channels (two front-end analog modules). Separating the functionality of the front-end digitizer card from the analog signal processing avoids mixing low-level analog and digital signals. The digitizer cards are located in crates shared with the calorimeter tower electronics placed inside the detector directly on the back of the calorimeter. These crates centralize common functions of clock, trigger synchronization, and slow-control processor. Each crate will process 8 front-end analog modules (16 after upgrade), and one for the Massless Gap front-end analog module. Initially, there will be a total of 416 digitizer cards. Each card sends its digitized data to a data-collection card.

Switched capacitor arrays (SCAs) will be used to store the analog data until digitization. The SCA is a 4K addressable sample-and-hold analog memory, divided into 16 channels of 256 cells each as described in Section 8.3.2. It samples and stores, in parallel, 16 analog inputs at the SSC crossing rate of 60 MHz. It can simultaneously store a 12-bit analog sample in one cell and retrieve another such analog sample from a different cell. The analog memory is controlled by a specialized integrated circuit, the Address List Processor (ALP), which manipulates stored data addresses (pointers) through a series of bookkeeping lists.

The architecture of the readout and digitization of the SCA (and of the readout of the digitized results) depends on how many bits per shower maximum detector strip are needed by the Level 2 trigger. If only hit strip signals (1 bit per channel) are needed, the required discriminated signals are available directly from the analog signal processing ASIC's, without digitizing the SCA samples. In this case, the SCA cards as designed for the calorimeter tower channels (see Section 8.3.2) will be used for the shower maximum detector readout. This readout technique was used as the baseline for the cost estimate. If pulse height information (*e.g.*, 8 bits) from the SMD is needed by the Level 2 trigger then the analog-to-digital conversion must be carried out just after the Level 1 trigger.

The Level 1 Accepts occur on average every 10 μ s. This time interval is sufficiently long to allow the use of a single ADC to convert three samples in each of 16 SCA data channels. The 16 channels are multiplexed to a common analog bus connected to a single inexpensive 8-bit 20-MHz Flash ADC. In order to preserve the full 12-bit dynamic range, a range-detection system composed of two comparators and three amplifiers (gains 1, 4, and 16) is placed in front of the FADC. The result of each conversion is a 12-bit number, which is stored in a digital FIFO buffer. An alternative solution will be to use a commercial monolithic or hybrid 12-bit 20-MHz Flash converter when it becomes available at a reasonable price. The operation of the FADC and the FIFO buffer is controlled by the ALP. The total time to process three adjacent crossing samples in each of 16 channels is 9 μ s.

Data collection boards

The data collection board is the last element of the SMD front-end system. It interfaces to the DAQ event builder. If SCA digitization occurs after Level 2 Accept, the SMD data collection board can be identical to the calorimeter tower data collector.

If SCA digitization occurs after Level 1 triggering, then the data collection boards must provide higher aggregate bandwidth, and they must provide data to the Level 2 trigger processor. In this case, the data collection boards are buffers common to the Level 2 trigger and DAQ systems. Each data collection board gathers data from one complete ϕ sector (barrel plus endcap). There will be 32 data collection boards, each with 24 inputs (14 barrel plus 10 endcap), located in four SDC standard DAQ crates with standard DAQ CPUs and trigger synchronization units. These crates could be located either inside the detector or in the surface-level counting house. A 10 Mbyte/s link with a simple synchronous protocol will be sufficient




FIG. 8-27. Block diagram of the shower maximum detector front-end electronics system.

to transfer all the data from 64 readout channel on a digitizer card to the data collection board. Some of the data e.g., 8 bits will be extracted from these boards and sent to the Level 2 trigger. Data will also be stored in the data collection boards pending the Level 2 decision and sent to the event builder by the standard DAQ path following a Level 2 Accept.

The average event size for the SMD is estimated to be about 30 kBytes, and a data throughput of 10 Mbyte/s per crate is adequate for a 1 kHz Level 2 Accept rate.

8.4. Straw tube front-end electronics

8.4.1. Design requirements

For a high precision drift tube or straw tracking system, it is necessary to accurately measure the time of arrival of the first electron (or cluster of electrons) at the anode. This, combined with the desire to operate with as low a gas gain as possible, implies the use of a low noise preamplifier with risetime sufficiently fast to provide the desired time resolution, but sufficiently slow to provide acceptable signal-to-noise.

In addition, because of the high rate of pulses on individual wires—for the inner wires at $\mathcal{L} \sim 10^{33} \text{ cm}^{-2}\text{sec}^{-1}$ the rate will approach 5–6 MHz—excellent double pulse resolution is very important.

The basic design goals are:

- $< 0.75 \text{ ns}$ time accuracy in order to ensure spatial precision of $< 150 \mu\text{m}$.
- Double pulse resolution of 20–30 ns. We have adopted the specific goal of having the return to baseline for a single cluster be less than 15 ns.
- Semi-gaussian shaping to minimize baseline shifts and noise from parallel current sources.
- On-chip Level 1 storage for 3–4 μs and Level 2 storage for a latency interval of order 50 μs .
- Ability to simultaneously acquire new data, store data, and read out interesting data without self interference or causing interference to other systems.
- Ability to withstand $\sim 1 \text{ MRad}$ and $> 10^{14} \text{ n/cm}^2$ over the life of the experiment.

8.4.2. Overview

The Straw Tube front-end electronics system can be divided into physical objects, logical blocks and data flow paths. The principal physical objects of the system include:

- Front-end Boards (FEBs) mounted close to the straw tubes. They receive signals from the straw-tubes, perform some immediate processing and forward the results to the crates.
- Crate Interfacing Cards contained within SDC standard crates (Section 8.9.4). They receive data from the FEBs and make it available to the DAQ system.
- Crates distributed within the SDC volume.

The system functions that support the operational requirements, are:

- High Voltage Connection and Distribution System—the means of providing power to the detectors and providing the interface between the Anode and Cathode and the active electronics section.
- Preamplification, Shaping, and Tail Cancellation—the initial analog processing steps serving to maximize the signal to noise (S/N) ratio and minimize the double pulse resolution time.
- Discrimination—the conversion from the leading edge of an analog pulse to an accurately timed digital pulse.
- Time measurement—the conversion from the leading edge of a discriminator pulse to a digitized time of occurrence.
- Level 1 Storage—the temporary storage of event data while the Global Level 1 Trigger System calculates whether a given crossing had possibly interesting data.

- Level 2 Storage—the temporary storage of event data from events accepted by the Level 1 Trigger while the Global Trigger System completes a more exhaustive calculation (from 10 to 50 μ s).
- Trigger Generation—the *local* generation of track segments from *local* discriminator signals; track segments so found are passed on to the Global Trigger System in time for use in the Level 1 trigger processing.
- Trigger Interface—receivers and buffers to distribute the minimum set of signals *from* the Global Trigger System - Clock, Level 1 Accept, Level 2 Accept, and Level 2 Strobe.
- Data Readout—buffers, drivers, and interfaces to the Data Acquisition System (DAQ).

The data flows of interest are:

- Data for events accepted by Level 2 to DAQ.
- Trigger Data (track segments) to the Global Trigger.
- Control and monitoring data and signals via the DAQ Control/Monitoring Network (Section 8.9.8).

A schematic diagram of the front end electronics is presented in Fig. 8-28. The preamplifier, shaping amplifier, and discriminator (ASD) will be implemented in a bipolar integrated circuit—robably with eight channels in a package. The time measurement, Level 1 storage, Level 2 storage, and readout control will be implemented in CMOS integrated circuits. The most conservative approach proposes the use of two chips, a Time Memory Cell (TMC) which includes the time measurement and Level 1 storage for four channels, and a Level 2 Buffer (L2B) which includes the Level 2 buffer and readout control logic. There are, however, significant advantages in terms of total die size and cross talk to implementing both functions on a single chip and this option will also be pursued. Prompt signals indicating the presence of stiff track segments are generated by Trigger chips which are also implemented in CMOS.

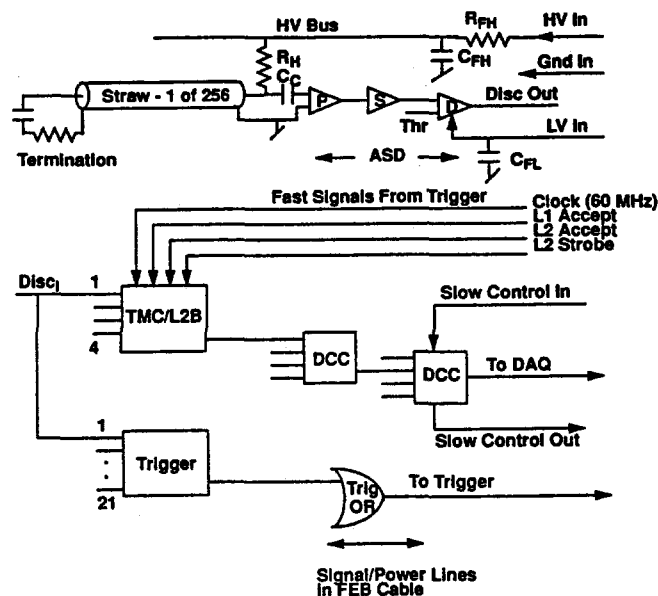


FIG. 8-28. Block diagram of the front-end electronics for the straw tube detector.

It is currently planned that the electronics will be built in Front End Boards (FEBs) of about 10 by 13 cm. The boards will be mounted close to the signal ends of groups of straw tubes, with each FEB including the time measurement for 159 stereo straws or the time measurement and trigger generation for a module of 212 trigger layer straws. The total number of FEBs in the system will be approximately 720.

The data from all channels on an FEB will be collected by one to five data collection chips (DCC) of which one, the Front End Readout Controller (FERC), will package the data and ship them to the crates mounted on the back of the calorimeter.

Each FEB for the axial (or trigger) wires will include 27 ASD chips, 14 Trigger chips, 53 TMC chips, 53 L2B chips, and ~ 5 DCC chips (plus passive components and, possibly, a few receiver/driver chips). Because of the large ratio of total silicon area to substrate area ($\sim 30\%$), the die will be attached by using either a) chip-on-board (via wire bonding, tab bonding, or flip chip technology), or b) advanced surface mount packages (~ 0.5 mm pin spacing) with packages mounted on both sides of the board. It is also likely that some of the chips, such as the TMC and L2B may be produced as multi-chip modules (MCM's) which are mounted perpendicular to the primary board.

8.4.3. Data flow and readout

It is expected that a single cable can carry the fast and slow signals and the power supply voltages between a single FEB and its associated crate, mounted on or near the back of the calorimeter. However, the possible advantages of optical fibers, especially for the transmission of fast signals, will be investigated.

It is presently planned that six cables from six FEBs will terminate on each of 120 Crate Interface Cards (CICs). These cards will be placed in eight tracking crates (four per barrel end). Each CIC will contain minimal logic for transferring to the DAQ system the event data originating on its six FEBs and for transferring the track segment data (for axial super layers) to the Trigger system. Because of the potentially very high data bandwidth from one tracking crate, the data flow will be split into four fiber optic data channels per crate. That is, there will be three or four CIC's feeding one fiber over a local private bus. The fiber interface will be the standard DAQ fiber adapter card.

An estimate of the data flow rates at various stages of the data stream has been made. These include overhead words to tag the data source, beam crossing, and trigger numbers. We find 1.5 Mbyte/s from each front-end board and 150 Mbyte/s from each front-end crate. The system aggregate data rate is 1.4 Gbyte/s.

Trigger path

Prompt signals from the axial trigger superlayers will be used to indicate the presence of stiff track segments. These trigger data will be matched with calorimeter, muon, and silicon tracker data by the trigger processor. Since the straw-tubes are arranged in radial *towers* pointing to the beam interaction point and half-cell staggered, two different types of trigger algorithm, time difference and time sum (mean timer), are possible.

The method of measuring the time difference within cells on a track is shown in Fig. 8-29. The time difference is inversely proportion to the transverse momentum p_t for properly arranged straws. By changing the clipping time, T_c , the p_t threshold value is adjustable over a range of a few GeV/c. Although the straw signal has a timing ambiguity that depends on the z-position of the track (dependent upon signal propagation delay down the straw and partially compensated for by η -dependent particle propagation delays), the ambiguity is $< \pm 3$ ns, and thus a time difference or mean timer circuit can be used for identifying the bunch crossing. There are several schemes that use combinations of this information to extract a high p_t track segment trigger signal.

The muon chambers and shower maximum detectors will have 1024 ϕ bins, while the outermost straw superlayer has about 2600 ϕ bins. Thus track trigger signals will be combined at the FEB in order to reduce the number of cables (Note that the required trigger bandwidth from the module, even with a limited number of ϕ bins, completely dominates all the other required bandwidths). Further study is required to optimize the trigger generation logic to be as effective, reliable, flexible, and inexpensive as possible.

Figure 8-30 shows the trigger-information flow from the straw-tube detector and indicates the return path of Clock, Level 1 and Level 2 signals from the Trigger System to the Straw Tube front-end system. A key element in both the time difference and mean timer circuits is a delay line. The Time Memory Cell (TMC—Section 8.4.8) is itself a combination of delay lines, memories and precise delay elements. It may be natural to include the trigger segment-finding circuits in the TMC, but it is also possible to implement the circuits in a separate chip. This is a complex question which involves the relative quantization of TMC and Trigger signal paths and basic packaging technologies. The initial plan is to have separate chips for segment finding.




FIG. 8-29. Time difference method of segment finding.




FIG. 8-30. Relationship of the straw tube front-end electronics to the Level 1 trigger processor.

It is possible that stiff track trigger signals processed with several different p_t thresholds will be needed. If so, transmitting the data via signals with multi-value levels (rather than strict binary levels) might permit minimizing the number of cables required to carry the data. More detailed p_t and track position information could be made available for the Level 2 trigger, but this would involve additional trigger data paths. For instance, it would be possible to send 500 bits of information to the Global Level 2 Processor in 5 μ s by using a 100 Mbit/s bit-serial line.

Control path

Each module or FEB has a serial network interface which is used for slow control of the system. All the monitoring, calibration, diagnostics are done through the serial network. The serial networks in each FEB are linked via the Crate DAQ interface to the SDC Control/Monitoring Network. Supervisory and diagnostic control is thereby available to any legitimate processor in the network. The FEB includes test pulse circuits for the preamp and the TMC. It will be possible to bench test any module or FEB via any system capable of connecting to the serial network interface. This serial network may be a standard network (such as JTAG/IEEE 1149.1) or may be SDC specific.

8.4.4. Analog signal processing

Each straw channel will require a preamplifier, shaper, discriminator, time converter, Level 1 and Level 2 storage arrays, and a data sparsification unit. Bipolar technology offers optimal performance for the preamplifier, shaper and discriminator. Bipolar's high gain-bandwidth at low power is unmatched by any other commercially available technology; it has the best noise performance for shaping times of less than 20 ns; and provides good matching of the control voltage between transistors. CMOS, because of its high density and potentially low power dissipation, is clearly the technology of choice for the time measurement block, which must include memory for the Level 1 and Level 2 storages, and must include much of the overall control logic.

Properties of the straw signal

Due to the high occupancy and low operating gain in the straw tube system, it is necessary to carefully consider the characteristics of the straw tube in designing the readout system.

The drift velocity for ionized electrons in CF_4 , the fast gas being explored for use in straw tubes, is about 100 μ m/ns. Achieving a position accuracy of 100 μ m therefore requires sub-nanosecond timing accuracy from the electronics. Since there is no plan to store charge information for off-line analysis, it will be important to trigger on the avalanche resulting from the initial drift electrons to get the most accurate timing information. This condition severely restricts the amount of charge available to trigger the electronics. Most of the signal from the straw is induced by the motion of positive ions towards the cathode, a process that takes about two hundred microseconds to complete. Much of this signal must be

truncated; for a typical straw tube only about 16% of the total charge for a single avalanche cluster has been collected after 6 ns. Extending the measurement time decreases the timing accuracy for signals of different amplitude due to time slewing. Decreasing the measurement time reduces the available signal and, in gain limited applications, may severely affect the signal to noise.

Simulations of the effects on the timing resolution of intrinsic noise from the straw and from the preamplifier with multipole shaping have been performed. These studies have led to the conclusion that a five to seven nanosecond shaping time gives reasonable signal to noise, without compromising the goal of a locally determined sub-nanosecond timing accuracy. This gives an approximate available signal of 0.64 fC (4000e) per drift electron for a terminated straw tube operated at a total gain of 5×10^4 . Measurements of prototype straw tubes have verified that a position resolution of $< 150 \mu\text{m}$ may be obtained with this peaking time.

We have developed and tested a lumped-sum model of the straw tube as a lossy transmission line. The determination of signal-to-noise for a straw tube with a transmission line termination requires careful analysis. The complex impedance of the lossy straw tube requires a complex rather than purely resistive termination. A suitable passive circuit is formed by adding a small capacitor in series with the termination resistor. The signal, of course, is absorbed at the termination end and therefore only half of the total charge is available to the readout amplifier. In addition, when formed with only passive components, the intrinsic noise in a measurement is dominated by the termination (assuming a well designed preamplifier). In this situation, the choice of symmetric shaping helps reduce the noise. While simulations show that significant additional noise reduction is possible by replacing the passive components at the termination end with an active load, at a power cost of about 3 mW per channel, the additional cost and detector design complexity probably argues against using this type of termination.

8.4.5. Bipolar preamp and shaper

Several different designs for bipolar preamplifiers and shapers have been pursued, at KEK using an advanced NTT process, and at Pennsylvania using both a complementary AT&T process and a high speed Tektronix process. The discussion below concentrates on the AT&T preamplifier design, but similar performance should be obtainable from the NTT design. The full multichannel Preamp/Shaper/Tail Cancellation/Discriminator chain will first be realized in the Tektronix process. It is important to note that the circuit design is, to first order, independent of technology and any of the above technologies could be used to fabricate the production devices.

Circuit design

For the preamplifier design, we have chosen a cascoded common-emitter configuration with a dominant pole primarily determined by the feedback network. A self-biasing circuit allows the input transistor to be large and to operate near its noise optimum, while requiring only a modest supply voltage of 3 to 4 volts. The power dissipation is therefore relatively low and the circuit may be easily implemented in advanced, high-speed technologies that characteristically have low breakdown voltages. The common-emitter structure is duplicated and forms a pseudo-differential input for the fully differential shaper.

A differential structure has been implemented in the preamplifier and shaper to help eliminate sensitivity to external sources of noise conducted through power supply lines or radiated into the chip itself. While this slightly increases the thermal noise compared to a single ended input, it is felt that it will aid significantly in noise rejection (such as from RF pickup in connections to the straw tube) and will provide a more robust system, at least for initial field tests.

The shaper has one zero, matched to cancel the dominant pole of the preamplifier, and three equivalent integrations that limit the bandwidth and maintain pulse symmetry. Good matching is achieved in the pole-zero cancellation by choosing components of the same types to dominate the pole and the zero. The integration poles are formed in the collector nodes of each differential pair. The resistance of the collector resistor times the stray capacitance at each node sets a time constant of about 1.7ns for each stage. These time constants match to within a few percent between stages, but are expected to vary more than that

chip-to-chip. Manufacturers' four-sigma guarantees are about $\pm 25\%$, but quoted *typical* numbers are less than $\pm 15\%$, and our experience over a small number of wafers is that the variation has always been less than $\pm 10\%$.

The intrinsic amplifier noise is primarily determined by the size of the input transistor and its quiescent current. In the design, the noise performance was optimized for an amplifier with a pure detector capacitance of 5 to 10 pF. The quiescent current was set to about 0.5mA and an input transistor of $75\mu\text{m}$ emitter length was used to achieve a base resistance of about 15Ω . We have compared the measured results and SPICE calculations using transistor models provided by AT&T for several values of detector capacitance. The agreement between measurements and model-based calculations encourages our reliance on SPICE-based modeling to predict the performance of system blocks. Independent noise measurements of the preamplifiers connected to straws confirm a noise performance similar to that predicted.

8.4.6. Bipolar discriminator

We are presently exploring the design of very low-power timing discriminators for use with the preamp and shaper. The present design goals are:

- Power — 5–10 mW
- Minimum Threshold — 10 mV
- Hysteresis — 2 mV
- Time Slewing — 1 ns (3–300 mV overdrive)
- Output Drive — 150 mV differential into 5 pF

Our designs attempt to take full advantage of the benefits of advanced bipolar technology. Since base-emitter matching between transistors is about 1 mV, it is possible to design for a reliable minimum input signal of 10 mV or less without input trimming. High unity-gain bandwidth transistors allow for the use of cascaded gain blocks with feedback to implement hysteresis and reduce time slewing. Drawing on previous experience in the design of differential discriminators, we have developed several possible configurations suited to particular bipolar technologies and are planning to implement a design in our next bipolar run. The near threshold performance of one design was modeled using Tektronix SHPi technology. From 3 mV to 12 mV, the largest overdrive used, the output delay shortens by about 1 ns. From about 500 μV to 3 mV the output delay changes by nearly 2 ns. With a power requirement of 7 mW this circuit should easily satisfy our design goals.

New prototypes

As noted above, a multi-channel (8 channels per chip) preamp - shaper - discriminator with detector tail cancellation has been designed and will be fabricated in the Tektronix SHPi process. The low stray capacitance of this process allows a much lower power dissipation than in previous designs, 7 mW in the preamp and shaper and 7 mW in the discriminator. Revision of the preamplifier to increase the open loop gain has reduced its sensitivity to the value of the detector capacitance. The input impedance has been tuned to be approximately 120Ω over the useful bandwidth of the preamplifier, with a smooth rolloff at high frequency. The tail cancellation circuit is based on a design by John Oliver (Harvard).

Radiation hardness

Six channels of the AT&T preamplifier have been exposed to Co^{60} radiation at BNL. Three channels were exposed to 1 Mrad and three to 2 Mrad. A small gain loss of approximately 5% was measured for those channels exposed to 2 Mrad and there was essentially no change in the noise. (The measurements actually indicated a small decrease, but this is probably due to systematic uncertainties in the measurements.) Rise-time measurements indicate an increase of approximately 0.6 ns for all channels and only one channel out of six had more than a 1 mV shift in the DC output voltages.

8.4.7. Drift time measurement

The Time Measurement block takes randomly occurring discriminator signals and provides a digital output that accurately defines the time for the front edge of each discriminator pulse associated with an event accepted by the Level 1 and Level 2 triggers. Two different devices have been under development to satisfy these functions—the TMC (Time Memory Cell) developed at KEK and the TVC/AMU (Time to Voltage Converter and Analog Memory Unit) developed at Pennsylvania. Both devices have demonstrated a multiple-hit time resolution of better than one nanosecond. At the moment, the TMC, which contains the Level 1 but not the Level 2 storage is in a more nearly final form and is used as the basis for this conceptual design. However, it is still necessary to develop an integrated circuit to store the output of the TMC during the Level 2 Trigger latency time and to provide an interface to the Data Collection Chip. The Level 2 Buffer chip (L2B) is an all-digital device, and is considered to be relatively straight-forward to develop. Possible interference from the high-speed digital signals from the L2B to the preamp inputs is a significant issue that must be addressed. It will also be necessary to increase the radiation hardness of the chip set. Some work on the TVC/AMU will continue and it should be available as an alternative device if some unforeseen problem should arise with the TMC.

8.4.8. TMC

The Time Memory Cell (TMC) chip is a low-power time-to-digital converter chip which includes the first level buffer inside the chip. The TMC records the history of the input signal in a memory array using a fully digital method. The input signal is fed to the data line of each CMOS memory cell, and each "write" signal to the memory cell is delayed column by column with a variable delay element which is controlled by a feedback circuit. Thus in the present chip, a 60 MHz write signal is delayed about one nanosecond per column across the 32 bits of column width; the write pointer is then advanced to the next row in the 32 row array. This yields a nanosecond by nanosecond history of the input signal for 1,024 nanoseconds (32 rows by 32 columns). Table 8-8 summarizes the specifications of the present chip (TMC1004).

Table 8-8
TMC1004 and TMC-SSC Specifications

	TMC1004	TMC-SSC
No. of Channels	4 channel	4 channel
Least Time Count	1 ns/bit	2 ns/bit
Time Range	1.024 μ s (4 ch), 2.048 μ s (2 ch) or 4.096 μ s (1 ch)	4 μ s
Clock Frequency	60 MHz	30 MHz
Time Resolution	$\sigma = 0.52$ ns	$\sigma = 0.75$ ns
Number of Rows	32	128
Number of Columns	32	16
Data Encoding	32 bit to 5+1 bit	16 bit to 4+1 bit
No. of Pins	I/O pins = 54 Power/Gnd pins = 34	I/O Pins \sim 50
Supply Voltage	3.0 V	3.0 V
Power Consumption	7 mW/ch	\sim 8 mW/ch
Chip size	5.0 \times 5.6 = 28 mm ²	6 \times 7 = 42 mm ²

The time resolution of the TMC1004 is $\sigma = 0.52$ ns (Fig. 8-31). This value can be explained as a combination of the digitization error ($\sigma_{dig} = 1/\sqrt{12} = 0.29$ ns) and the TMC inherent error ($\sigma_{TMC} = 0.43$ ns).




FIG. 8-31. Measured time resolution of the TMC.

To increase the buffer length while keeping the Si area within an acceptable level, we propose to increase the least count to 2 ns/bit from the 1 ns/bit in the present design. If the σ_{TMC} of new design (the TMC-SSC) is the same as the present one, the time resolution will be

$$0.75 \text{ ns} = \sqrt{(2\text{ns}/\sqrt{12})^2 + (0.43 \text{ ns})^2},$$

which is expected to be within the requirement of the wire chamber readout. For the TMC-SSC we would then have 16 columns (at 2 ns each) with 128 rows (to give a 4 μs Level 1 delay time) driven by the same 30 MHz clock. The increase in Si area is only 50% even with the quadrupled buffer length, because the control logic and interconnection pads occupy about 70% of area in the present chip.

The specifications proposed for the TMC-SSC design are shown in the second column of Table 8-8. A block diagram is in Fig. 8-32. In the TMC-SSC, we propose to encode 16-bit data to four data bits plus one carry bit, and to use a system clock of 30 MHz (32 ns period). As described in the next section, it is still possible to extract event data synchronized to a trigger signal of 16 ns resolution. The encoding scheme reduces both the number of output pins on the integrated circuit package as well as the size of the data packet. Use of the lower clock frequency results in an easier chip design problem and reduced power consumption.

8.4.9. Level 2 buffer

Figure 8-33 shows a block diagram of the Level 2 Buffer (L2B). The L2B consists of encoder logic, a buffer memory, a buffer controller, a bus interface, TMC control logic, and trigger control logic. The buffer controller has two pointers; the buffer operates as a ring buffer. This chip is a fully digital chip using ASICs in the form of gate arrays or standard cells. In the L2B, the data from the TMC are reconstructed to drift time as shown in Fig. 8-34. Since the maximum drift time of a straw is around 30 ns, several rows of data have to be read out from the TMC for each trigger (to allow for the uncertainty in beam crossing). The L2B reconstructs the drift time from these data. Referring to the phase synchronization of the Level 1 trigger signal, the L2B recognizes the starting point of the data and then calculates the drift time. Although the figure shows only one data point, it is possible to process multi-hit data.




FIG. 8-32. TMC-SSC block diagram.




FIG. 8-33. Level 2 buffer block diagram.




FIG. 8-34. TMC-SSC and L2B data encoding scheme.

8.4.10. Electrical interference

One of the most difficult problems to model accurately is the potential interference (cross talk) from one part of the system to another or to another detector system. It must be remembered that the basic sensitivity of the front end electronics is about 1 fC for a signal with a pass band centered at about 60 MHz

(~ 5 ns t_r , t_f), and that each electronics channel is connected to a detector that is 3 to 4 meters in length. Thus, in the time domain, a single 5 V transition with at $t_r < 5$ ns and a signal to preamp (or straw) capacitive coupling of about 200 aF (2×10^{-16} F) is sufficient to cause a discriminator to fire. In the frequency domain, the 3 to 4 meter straw detectors act as dipole antennas tuned to a $1/4\lambda$ frequency of about 60 MHz—distressingly identical to the ~ 60 MHz beam clock.

Since the actual calculation of most of these interactions is beyond the present state of the art, it will be necessary to be as conservative as possible in designing the electronics to avoid oscillatory behavior. In addition it will be necessary to fabricate relatively large-scale test detectors that will simulate many of the final relationships just to check for unfortunate interactive behaviour.

There are a variety of strategies available to implement *conservative* designs, but they are variations of two basic themes - avoid causing problems by keeping signals small, slow, and differential, and avoid seeing problems by keeping inputs shielded, using differential structures to subtract common mode noise, and keep the potential signal paths short (avoiding *ground loops*). These strategies are all fairly obvious and straightforward, but only large scale tests will give us confidence that whatever measures we adopt are sufficient.

The arguments for the various strategies are:

- Small signal size—since the energy transferred via capacitive coupling is proportional to the ΔV , using low level (perhaps a few hundred mV) signals can reduce the capacitive coupling effects by nearly a factor of 25.
- Slow signals—since the amplifier bandwidth is limited on the low side and since inductive coupling depends upon di/dt , a slow edge ($t_r, t_f \gg 100$ ns) will couple very inefficiently and a reduction of cross talk of one to two orders of magnitude is possible.
- Differential Signals—even fast, large edges, if they are generated and transmitted differentially, will be less likely to cause crosstalk. The amount of crosstalk reduction varies with the exact means of driving and transporting signals, but is acknowledged to be generally in the region of a factor of ten.
- Shielding—obviously by reducing capacitive and inductive linking of circuits, it is possible to control cross talk. Very good industrial strength shielding can provide 60 to 90 db of rejection over frequencies from tens of Hz to GHz - unfortunately this involves unacceptable amounts of copper and steel (for the magnetic component of the field, \vec{B}). Nevertheless, the straw cathodes (uniformity, electrical depth, and termination) can have a strong effect on the ability of the system to avoid oscillation.
- Differential Structures—at the amplifier input can, if properly connected to the detectors, subtract out a great deal of the system wide *common mode* noise. The prototype preamplifiers are designed with this in mind.
- Short Signal Paths—(and matched differential signal paths) tend to make inputs less effective antennas, and also make outputs less effective radiators of signal.

Some of these strategies will affect how we partition and package the various functions of the straw electronics system. For example, the prototype TMC involves a relatively large number (16) of single-ended 5 V output lines with potential transitions every 32 ns. These signals go directly into the L2B. While it would be possible to go to fully differential low level signals (as we will probably have to do for the connections from P/S/D to TMC), it may be more efficient and effective to repartition the functionality and include the L2B functions within the TMC chip (perhaps two channels per die). This would eliminate many fast single-ended signals and would cut down on the total number of pads needed on the die, but would force the redesign of a complex and tightly packed system.

In any event, a series of tests using the eight channel ASD at full density (*i.e.*, about 200 channels on a few hundred square cm) on various prototype chambers will be done in 1992 to try to understand the severity of the problem and to provide some partially quantitative levels. The results of these tests will have to be fed back into both the electrical and physical design of the system.

8.5. Muon front-end electronics

The SDC detector muon system covers more than 2500 m^2 with $\sim 90,000$ wire chambers, 2240 scintillators with phototubes on two ends, and 2256 scintillators with a single phototube. The geometry of the muon system determines the modularity and connectivity of the electronics.

The barrel region consists of three concentric octagonal barrels—BW1, BW2, and BW3—covering $|\eta| < 1.0$. BW1 is inside the toroid. BW2 contains the single layer of scintillators, $60\theta \times 4\phi$ scintillators each with two phototubes. Each barrel layer has θ measuring chambers, BW1 and BW3 also have ϕ measuring chambers, and BW3 has stereo drift cells.

The wire chambers and scintillators covering the outer radius of the endcap regions form the intermediate muon system. They are also segmented into ϕ octants containing θ , ϕ , and stereo chambers. Some of the drift tubes run transverse to the beam direction along the face of the toroid to extend the η coverage to 1.6. Per endcap octant, there are $10\theta \times 4\phi$ scintillators with phototube readout on both ends.

The very forward portion of the muon endcaps continues the η coverage to 2.5. In this area the tubes are of smaller diameter but remain projective. There are θ chambers but no ϕ chambers; the orthogonal coordinate measurement is provided by stereo wires since ϕ oriented tubes are not practical at these small radii. Per endcap octant, there are two layers of ~ 140 scintillators with single phototube readout. The θ segmentation is 33 and the ϕ segmentation varies from 3 to 2 in going from the outer to inner radius.

In each layer in all angular regions, projective θ pairs of wires measure the bend angle resulting from the toroid. The wires project to the interaction point. This facilitates the use of simple differential timing electronics for selection of high transverse momentum tracks at the earliest trigger level. Since the muon chambers are outside the return path of the solenoidal field, all tracks originating from the interaction point are radial except for multiple scattering. A selection of near radial tracks in the ϕ chambers rejects cosmic rays and tracks that have scattered significantly (soft tracks).

8.5.1. Front-end electronics location

The natural modularity of the muon system derives from its octant and z segmentation. In the barrel region each octant is divided in z into towers: 3 for BW1 and 5 for BW2 and BW3. Each endcap octant is a single pie-shaped unit. Since the physical extent of the system is so large, it is necessary to have clusters of electronics distributed around the detector. These clusters are called regional crates. Additionally, the trigger is formed by processing data from the outer two layers of the detector; hence the clustering of electronics must conform to the projective geometry of the detector. The regional crate count is $8 \times 5 = 40$ for the barrel and $2 \times 12 = 24$ for the endcaps, where the endcaps include the intermediate and forward regions.

The small signals from the chambers cannot tolerate cables as long as would be required to directly connect to the regional crates. To buffer these signals, minicrates of 192 wire chamber channels each provide the needed amplification and discrimination required to convert the chamber signals to differential digital pulses. The cables that connect these minicrates to the regional crates also provide outgoing power, thresholds, and test features.

The situation is similar for the scintillator phototubes except that in this case the minicrate is replaced by a single channel of electronics contained in the phototube base. The cable connecting the base to the regional crate provides power, a discriminator threshold, and high voltage monitoring. The testing of phototube channels is provided by fibers that direct laser light to the scintillation material directly.

The scintillator and wire signals arrive at the regional crate as differential digital pulses. These signals are processed by different cards for the two detector types—wire cards and scintillator cards. Both cards provide for Level 1 and Level 2 data storage, DAQ sequencing, and trigger processing. Wire cards contain test pulse distribution circuitry as well. A single scintillator card is anticipated for each regional crate. Between 18 and 32 wire cards are required for each crate with the forward endcap crates having the larger card counts. All wire cards will be of the same design although various portions of the trigger logic will be

disabled depending on the wire type— θ , ϕ , or stereo—to which the card connects. The remaining part of the regional crate is dedicated to the DAQ CPU, the Level 1 Clock Control Board, a test pulse fan-out card, and a board which provides trigger sequencing functions.

8.5.2. Wire chamber front end

The drift tube anode wires collect ionization from approximately 2 cm of track length, and the expected gas gain is the order of 10^5 . Threshold levels should be in the range of a few fC, and noise levels should be about 4000 electrons (1/2 fC). Risetimes of approximately 6 ns are adequate for spatial resolution, and longer risetimes are being studied. The sense wire resistance is as large as $30 \Omega/\text{m}$ (for 50 μm tungsten). The characteristic impedance of the tubes is 360Ω . For a 9 m tube, far-end termination is required to avoid unwanted signal reflections.

The front-end minicrates provide the circuits to amplify, shape, discriminate, and drive the individual chamber signals. In addition, these crates provide for injection of test signals at the preamplifier inputs.

The front-end electronics configuration is shown in Fig. 8-35. A small printed circuit board which contains a high voltage bypass capacitor and a pulse transformer is mounted to the end of each tube. This board also contains the high voltage connection for anode and field shaping. An RC terminator is attached to the far end of each tube. Signals from the tube are coupled via the pulse transformer to shielded twisted pair cable. A 2–3 m cable connects the tube to the the minicrate ASD card. Each ASD card accepts two 16-channel cables.

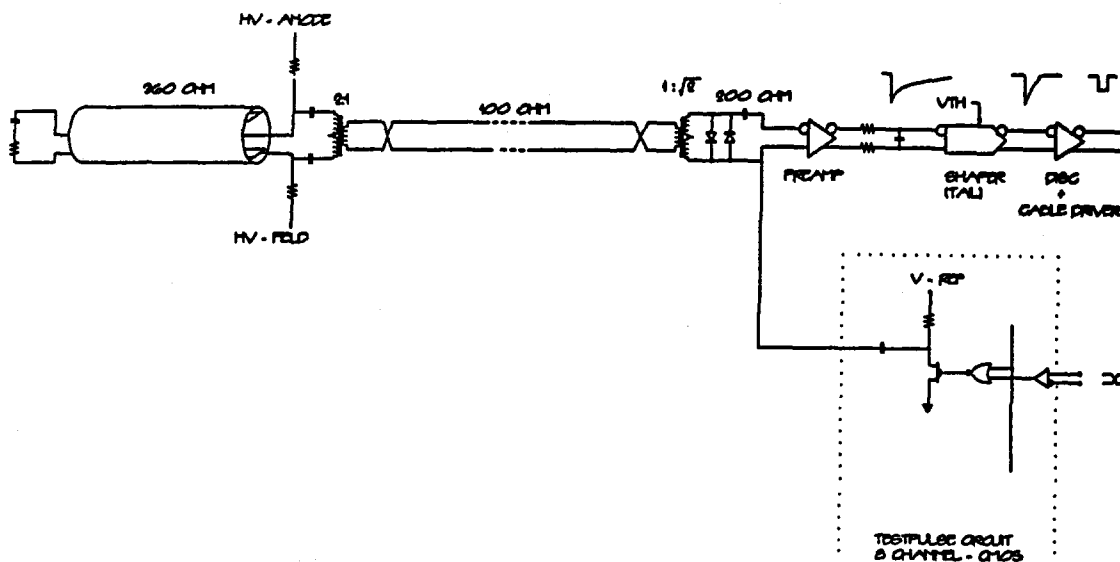


FIG. 8-35. Front-end electronics configuration for a muon drift tube.

The electronics for the input signal processing makes use of the 8-channel monolithic amplifier-shaper-discriminator chip (ASD) being developed for the straw tube tracker (see Section 8.4). Two other CMOS integrated circuits are being developed. The first device is a low-power differential cable driver to drive the discriminator signals from the minicrates to the regional crates. The second will be a device for injecting test pulse charges into the amplifier inputs while maintaining the channel-to-channel timing accuracy required for *in situ* testing of the Level 1 trigger logic.

Since the Level 1 trigger is based on time differences of the signals from the projective wires, a minimum of two accurately timed pulses is required to test the response of each trigger circuit. The timed pulses can be shared by all cards in the minicrate provided each 16-channel unit receives an enable line that controls its acceptance of the pulses. A common pulse bus will be connected to each minicrate from the regional crate. A single enable line will pass along each 16-channel cable and control the use of the timed pulses

by those channels. On receipt of the timing pulse, charge is injected into the enabled channels by an FET switch.

8.5.3. The phototube front end

The phototube base contains a single channel of high-speed comparator, its output driver, a high voltage generator circuit, and a high voltage monitor circuit. Several designs exist for bases of this type. A single cable connects this base to the regional crate. Figure 8-36 illustrates a base of this type.

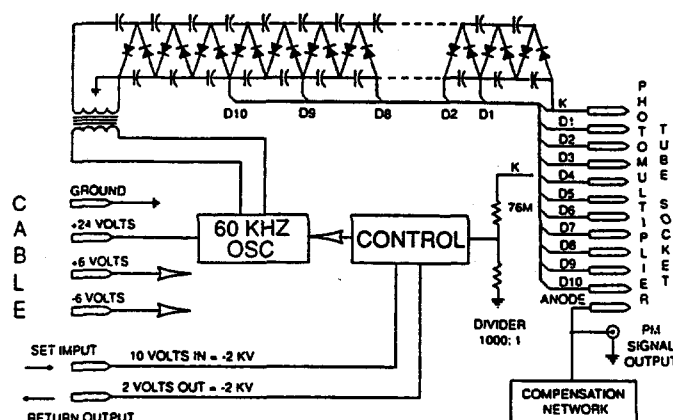


FIG. 8-36. Schematic diagram of the PMT base with Cockcroft-Walton HV multiplier.

8.5.4. Data storage and data paths

The regional crates receive the discriminated scintillator and wire signals in cards that provide the timing measurement, Level 1 and Level 2 data storage, and readout sequencing. The TDCs and data storage are implemented with the TMC chip used in the straw tube tracker (see Section 8.4 for details). The differential signals from the ASD minicrates first enter receiver chips that provide buffering and signal level conversion and are then stored in the TMC circuits, which provide time measurement and Level 1 buffering. The Level 2 buffering and readout for the TMC's are provided by the L2B chips, as in the straw tracker readout. Output of the L2B's passes directly into the board data collection buffers, onto the regional crate data path (currently considered to be VMEbus), and into the local DAQ CPU located in the regional crate. At this point the data flow conforms to the SDC DAQ standard.

The scintillator and wire signals are also sent to the trigger path, which is separate from the Level 1 and 2 data storage. For the scintillators this path contains the mean timers that determine the crossing time. For the projective pairs of wires the equivalent trigger path leads to the mean timers that apply the p_t determinations. The results of these trigger selections are also stored for later DAQ readout should the event pass the Level 1 and 2 trigger conditions.

A scintillator card receives up to 96 phototube signals. This card contains 24 TMC chips that digitize the time information and 48 mean timer circuits to provide beam crossing timing to the wire trigger circuitry. The cards provide registers for enabling individual channels and trigger processing.

A wire card receives 64 drift tube signals. This card contains 16 TMC chips for storage and the mean timers for drift time windowing. Circuitry to associate scintillator hits with chamber hits and to link projective chamber pairs in the two trigger layers is also present on this card. The card provides registers for enabling individual channels and trigger processing. Thus identical cards can be used for θ , ϕ , and stereo wires. The wire cards contain the circuitry for the distribution of test pulses to the amplifier inputs in the minicrates. Each 16-channel group of wires has a unique test pulse enable bit. The test pulse is generated on a test pulse fanout card and distributed separately to each minicrate.

8.5.5. Trigger overview

Inputs to the trigger are the projective θ wires, the radial ϕ wires, and the scintillators. The scintillator outputs provide the bunch crossing time for the trigger and the projective wires provide the momentum selectivity. For the barrel and intermediate regions, each scintillator is viewed by two phototubes, one at each end. For the forward region, there are two counters with single phototubes that are projectively aligned. A coincidence of two phototubes is used to suppress the singles rate. To accomplish this coincidence and to overcome the time variability due to the flight path differences within the scintillator a digital mean timer is used. The output of this circuit is the time average from the two phototube pulses. Figure 7-18 shows the cosmic ray test stand results for the prototype mean timer. The precision of the timer is far greater than required for this application. Additional position resolution along the length of the scintillators is also obtained from the mean timers and is used to help associate hits in θ and ϕ chambers.

For the Level 1 trigger, the projective θ wires select high p_t tracks as those that have small differences in pulse arrival times at pairs of projective wires. Time differences less than 80 ns to 320 ns are typical selection criteria. A digital mean timer provides a simple deadtimeless way to implement this function. A single timer can be configured to give multiple acceptance windows so that multiple p_t thresholds can be incorporated within the basic circuit with little complication. As an example, to find a muon in the barrel region at a given crossing time and above some p_t threshold, a track must be sensed in the two phototubes of a single scintillator, have the derived mean time occur at the proper the crossing time, and have at least one projective pair in each of BW2 and BW3 with a time difference that satisfies the p_t threshold. While left-right ambiguity crossovers can potentially inject low momentum tracks into this selection, the probability of two such crossover faults is rare. The position along the scintillator derived from the mean timer also localizes the ϕ region to search for the ϕ wire hits associated with the θ hits, providing the Level 2 trigger with four azimuthal regions.

Clearly, with many such coincidences possible from independently operated parallel circuits, there can be multiple muon signals from the same crossing, and coincidentally from different crossings. Algorithms for selection among the several muon candidates will have to be developed with the aid of simulations. In general, preference will be given to multiple muons from a single crossing over muons in consecutive crossings. Since signatures for multiple muons are known, and the probability for consecutive crossings to yield muons in the same region of the detector is vanishingly small.

8.5.6. Trigger implementation

Figure 8-37 illustrates a block diagram of the trigger. The scintillator card uses the mean timer to eliminate the track position dependent time variation of the signals. The mean timers have 4 outputs that can be used to subdivide a counter into 4 sections. For example, a barrel scintillator of 50 cm by 200 cm is divided into four 50 cm by 50 cm sections.

The treatment of the scintillators after this point becomes somewhat region dependent. The treatment of the barrel scintillators will be presented in detail for illustration. After time division, the scintillators in each muon tower scintillators form a 12 θ by 16 ϕ matrix. Since the scintillator-wire matching will be done separately in θ and ϕ , this matrix can be reduced to 12 θ and 16 ϕ ORs. Before these scintillator strip ORs are sent to the wire cards, they must be delayed to permit the maximum drift time to pass. This delay is metered in 16 ns ticks by running each OR through a shift register. The shift register output drives the backplane scintillation bus.

A wire card receives 64 discriminated chamber signals. In the trigger section, the wires are processed as projective pairs into five bits of p_t threshold for θ pairs and into bits of non-radial angle for the ϕ pairs. These pulses are widened to a time corresponding to the maximum drift interval and then compared to the scintillation bus signals which have been delayed for the maximum drift time. For ϕ wires a match with the appropriate scintillator strip results in the storage of the ϕ wire pair number for sequencing to the trigger arbiter (see below). For θ wires, a match includes linking between the two outer layers and a coincidence with a scintillation bus signal. When a θ match is detected, the highest threshold exceeded is output along

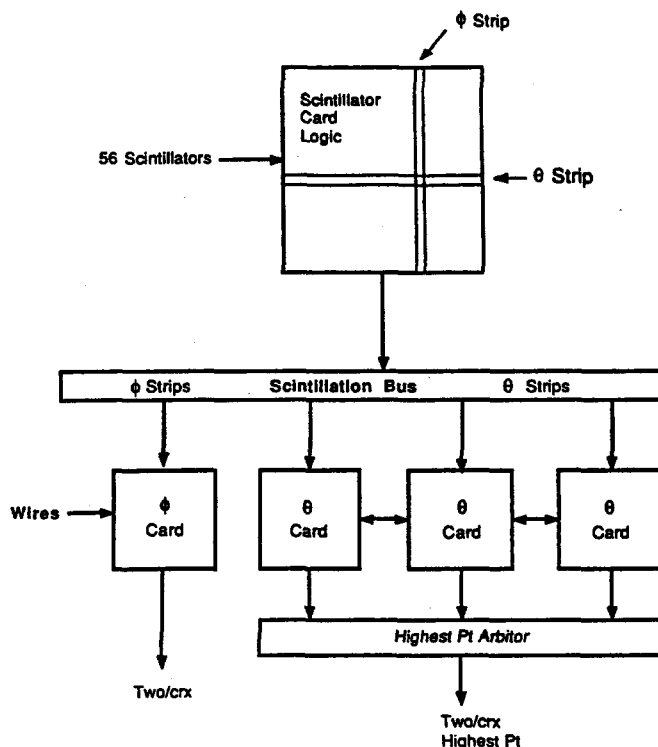


FIG. 8-37. Block diagram of the muon trigger system

with the ϕ strip address. This later provides some ϕ information even in the case when no ϕ wire pair is detected. All matching and timing parameters are computer programmable.

Inner layer θ and ϕ tube pairs can also be included in the trigger. Each inner pair must match with the appropriate scintillator strips and with a range of tube pairs in the outer layers. The range depends on the p_t threshold of the projective pair. This matching will be computer programmable for flexibility. The association of inner and outer layer pairs must span the individual card and regional boundaries. To accomplish this boundary-independent association, 16-pair signals from each card are transmitted to the adjacent card. This interconnection ties the edge cards of regional crates together via drivers and receivers.

A trigger item arbiter will receive as input the muon candidate hits and data from each of the wire cards. For Level 1, the arbiters transmit the two highest p_t muon candidates in η - ϕ bins defined by the scintillators: $10 \eta \times 32 \phi$ bins in each barrel tower and $8 \eta \times 32 \phi$ bins per endcap tower. For each track, a 2-bit code is sent where 00 represents no candidate and the other 3 combinations represent tracks above progressively higher p_t thresholds. For Level 2 data, the trigger item arbiters use the full p_t precision of the muon trigger logic and associate θ and ϕ tracks. The result is a 5-bit p_t for $10 \eta \times 1024 \phi$ bins in the barrel and $8 \eta \times 1024 \phi$ bins per endcap.

Figure 7-2 shows a simulation of a trigger formed with a single scintillation layer and two or more projective θ wire pairs for various drift time difference windows. The time differences are presented as spacial separations to 2, 5, and 10 mm corresponding to drift time windows of 60, 100, and 200 ns. In this plot it is required that both BW2 and BW3 contribute a pair to this coincidence. This concept forms the basis for the projective wire p_t selection.

8.6. Scintillating fiber option front-end electronics

8.6.1. Introduction

The purview of the front-end electronics system for the Scintillating Fiber Tracker extends from the outputs of the Visible Light Photon Counters (VLPCs) to the interfaces with the Trigger and Data Acquisition Systems.

The purposes of the system are:

- To detect and digitize the analog signals from VLPCs that denote the passage of particles through the scintillating fibers with a time resolution concomitant with the SSC beam crossing interval.
- To logically combine the digitized signals to detect track segments within individual superlayers.
- To logically link the track segments into tracks that traverse the four superlayers.
- To deliver to the Trigger system, each beam crossing interval, data on the linked tracks.
- To interact with the Level 1 and 2 trigger protocol and deliver fiber hit data to the Data Acquisition system.

The details of the scintillating fiber detector, the VLPCs and their cryostats are covered in Section 4.8.

A key feature of the Scintillating Fiber Tracker front-end electronics system is that all design requirements are implemented in technology that is now readily available. The design does not depend on future technological developments, although certain such developments may be incorporated in the design of the production version if appropriate.

A second key feature is that the system is almost entirely digital. Although the VLPCs generate analog signals, they travel only a few inches within the well-shielded cassette before being converted to digital signals. Thus, problems arising from ground loops and other noise sources are minimized.

The digital nature makes the system easy to design and easy to debug. In debugging, for example, the digital pipelines may be single stepped or operated at lower than normal clock rates to aid in diagnostics.

8.6.2. Physical overview

A block diagram of the front-end electronics for the Scintillating Fiber detector is shown in Fig. 8-38. The major system components of the system are:

- Amplifier-Shaper-Discriminators (ASDs). The ASDs amplify and shape the signals generated by the VLPCs and develop digital pulses for those VLPC signals that exceed a threshold.
- Trigger and Readout boards. These boards process the digital pulses from the ASDs to develop trigger data and event data. They also provide VLPC support—VLPC bias voltages, ASD power, and temperature monitoring.
- Interfaces to the Trigger and Data Acquisition (DAQ) systems through which data are transmitted to these two systems.

The VLPCs require an operating temperature of about 6–7 °K. They are installed within cassettes, which in turn are placed in cryostats. The complete system has 1024 cassettes, each containing 512 signal channels. The 512 signals from each cassette are carried over eight multi-conductor ribbon cables to its Trigger and Readout board, housed in a crate in a room temperature environment. The crates are distributed on the back of the calorimeter. Each of the eight ribbon cables transports signals from 64 channels. In addition, a 128-conductor ribbon cable carries power, control and monitoring signals from the readout board to the cassette.

Each crate can contain up to 16 Trigger and Readout boards together with the DAQ CPU and the Level 1 Clock Control Board. Communication among the boards is via the crate backplane. Signals to and from the trigger and DAQ systems are transmitted via fiber-optic cables using standard SDC protocol specifications.




FIG. 8-38. Overall block diagram of the Scintillating Fiber Front-End Electronics.

8.6.3. Amplifier-shaper-discriminators

Particles passing through the scintillating fibers typically generate <10 photons. These are converted to electrical signals by the VLPCs. The VLPCs have quantum efficiencies of up to 85% for light of 560 nm wavelength, and an electrical gain of 10,000 or more. They therefore deliver a pulse having a charge of a few times 10,000 electrons for each detected photon.

The Amplifier-Shaper-Discriminators (ASDs) are designed to deliver standardized digital pulses whenever the VLPC generates an output pulse above the threshold of the discriminator section of the ASD.

The basic requirements of the ASD are:

- An input signal/noise level of 4:1 or better for single photon pulses, i.e., $< 1/4$ photon equivalent noise.
- Pulse shaping such that the shaped signal returns to baseline in 16 ns, with a small time slewing over a range of amplitudes, achieving thereby single crossing resolution.
- Discriminate the shaped signals against a fixed threshold and generate short (<16 ns) digital pulses denoting "hits."
- Ability to drive terminated ribbon cable transmission lines to transmit the digital pulses to the readout crates.

The ASDs will be located within the cassettes. The advantages of this location include:

- A significant advantage with respect to their input noise level, since they will be much closer to the signal source.
- The signals carried on the ribbon cables between the cassettes and crates will be digital, rather than low-level analog.
- The ribbon cables may be of less expensive design—if the ASDs were to be in the crates, the ribbon cables would necessarily be of low-capacitance construction.

Preamplifiers

Several candidate preamplifiers have been tested with the VLPCs [46]. These include the QPA02 amplifier designed at Fermilab for the readout of silicon micro-strip detectors [47], the VTX amplifier designed at Fermilab for the readout of vertex tracking chambers for the CDF detector [48], and the TVC amplifier developed for the readout of straw tubes and muon chambers for SDC [49].

For reference purposes, the performance of these devices was compared to a transimpedance amplifier (TIA) designed and refined by Rockwell and Notre Dame specifically for use with VLPCs. Pulse-height spectra recorded for TIA, VTX, TVC, and QPA02 amplifiers have been shown [7]. In each spectrum, a series of individual peaks was clearly visible, corresponding to the simultaneous detection of integral numbers of photoelectrons. Clear peaks were visible for up to 6 photoelectrons in all cases because of the excellent gain dispersion characteristics of the VLPC, even when the signal is amplified with the fast preamplifiers.

Status of ASD work

The VTX preamplifier has been shown to work well with the current VLPC cassette. ASD cards have been designed and built at Fermilab for use with the VTX preamplifiers. They should be suitable for a test of 1000-2000 channels of Scintillating Fibers.

On the cards, the input analog signal is amplified by a Type 592 video amplifier with an adjustable voltage gain. Gains of over 100 are available. In addition, there is a pole-zero filter and two stages of integration for signal shaping. Following the shaper, a LeCroy MVL407 is used as a discriminator to generate a differential ECL output signal. Differential analog output signals are also available. With minor modifications, these ASD boards have been used in CDF since 1985.

8.6.4. Trigger and Readout board

A block diagram of the Trigger and Readout board is given in Fig. 8-39. These boards contain all the digital logic for performing three major functions: processing hit data (from the ASDs) and outputting pipelined information on linked tracks for the trigger system; providing Level 1 and Level 2 data storage for event data; and outputting event data to the DAQ system. Power dissipation is ~ 45 W.

These functions are implemented on two Application Specific Integrated Circuits, ASIC-1 and ASIC-2. ASIC-1, shown in Fig. 8-40, provides all Level 1 and Level 2 pipelined data storage, performs segment finding (i.e., clustering of fiber hits) within a superlayer, and forwards selected event hit data from 64 fibers toward the DAQ interface. ASIC-2 is primarily concerned with track linking.

The trigger data processing is done in two steps. First, data from each superlayer are examined for track segments by ASIC-1. A track segment is detected when a fiber from each of the four layers of the superlayer is hit. Only track segments representing a certain range of p_t are accepted. Secondly, track segment data from superlayers are linked by ASIC-2 to detect tracks that traverse the entire scintillating fiber tracker.

Neighbor communication

As shown in the block diagram, each board accepts the digitized signals from 64 ASDs, which represent 64 fibers. However, at the edges of the blocks of 64 fibers, tracks may be partially in a given block and partially in its neighboring block. To find all track segments without any gaps at block boundaries, it is therefore necessary to have hit data from five neighbor fibers also available. The local data and neighbor data must be synchronized before processing. To accommodate the time necessary to transmit and receive hit signals to and from neighbors, a two-stage latch is provided at the input of the board. Data are moved through the two stages at a 60 MHz rate. Of the 64 bits latched into the first stage, five bits are transmitted to neighboring boards. At the second stage, five bits from neighbors are latched along with the 64 local bits.




FIG. 8-39. Physical layout of the Trigger and Readout board. A board with the number of ASICs shown will handle 1024 fibers. For smaller numbers of fibers per board, the number of ASICs actually loaded would be less than shown.

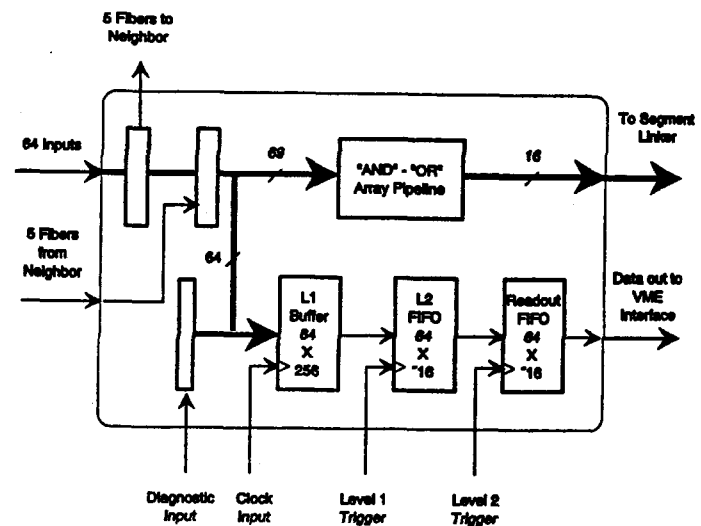


FIG. 8-40. A block diagram of ASIC-1, which implements the functions of Level 1, Level 2 and DAQ storage and track segment finding.

At the output of the second stage, the 67 bits are presented to the trigger logic; the 64 local bits are transferred to the Level 1 storage array which is implemented as a FIFO 64 bits wide and 256 beam crossings deep—sufficient for a $4\mu\text{s}$ Level 1 trigger delay.

Neighbor communication is also required for the track-linking operation. Here, the number of communicating neighbors is determined by the Barrel configuration and the p_t threshold. The number will never exceed 300 per Trigger and Readout board.

Track segment finding logic - ASIC-1

During each 60 MHz clock cycle, the AND/OR array pipeline accepts 69 bits of hit data, and outputs 16 bits of track segment data. The track segment data are those resulting from hit data of several clock cycles earlier.

The algorithm for track segment finding uses a fixed pattern of combinatorial AND/OR logic to find high p_t hits in a given superlayer with a resolution of one fiber diameter. This algorithm has been discussed in detail elsewhere [7]. To find all track segments of 10 GeV/c or greater momentum, 14 3-input AND gates and one 14-input OR gate are required for each fiber diameter position in the fiber array. If a hit is required to be found in every one of the four layers, a total of about 100 AND/OR gates is needed.

The algorithm is implemented in integrated circuit ASIC-1 that is custom-designed for this purpose. Each track-segment finding ASIC covers 16 fiber diameters of the four-layer array (64 fibers total) and therefore requires 224 3-input AND gates and 16 14-input OR gates. The algorithm can easily be implemented in currently available ASICs.

Each Trigger and Readout board contains 16 ASIC-1 integrated circuits.

Time slewing

In the present configuration, the fibers in each layer are approximately 4 m long, with the VLPCs located at the end farthest from the beam crossing point. The propagation delay (for light) from one end of the fiber to the other is therefore about 25 ns, greater than a beam crossing period. The slewing that would result from this is partially compensated by the difference in times of flight of the particles; the time of flight is longer for particles crossing the fiber nearer the VLPC. With this automatic compensation, the slewing owing to the point on the fiber where the particle crosses is about 12 ns. By keeping the sampling time at the input to ASIC-1 small—2 ns for example—the possible confusion in relating VLPC signals to the proper beam crossing number can be eliminated.

Track linking - ASIC-2

The segment-linking integrated circuit ASIC-2 serves to correlate track segments from three to five fiber superlayers (which are labeled B1 ... B6 from inner to outer) within a narrow angular cone corresponding to that expected for a track of transverse momentum 10 GeV/c or greater. For coverage to $\eta < 2.3$, the trigger is provided by the three inner superlayers, B3, B2, and B1.

Each Trigger and Readout board has 56 fiber-diameter positions for B3, 48 positions for B2, and 40 positions for B1. The set of 144 bits representing these 144 positions are processed by the segment linking ASIC-2s. Relative to a hit in layer B1, a 10 GeV/c track would produce a hit in layer B2 offset by 4 fiber diameters and a hit in layer B3 offset by 8 fiber diameters [7]. By using the proper set of combinatorial logic the entire set of fiber-diameter positions can be examined for tracks that satisfy a predetermined stiffness requirement.

Requiring segments from 3 superlayers serves to reject false tracks. It is also possible to make a range of p_t triggers by requiring that the tracks are less inclined than a 10 GeV/c track. For higher p_t thresholds the outer layers B5 and B4 could also be examined to aid in the determination of p_t at the trigger level.

Each of the six segment linker ASIC-2s on each Trigger and Readout board will cover 8 fiber positions at B1. This requires 12 fiber positions at B2 plus 4 fiber positions at B2 from neighbors on both sides, and 16 fiber positions at B3 with 6 from neighbors on each side. We would also include 20 fiber positions from B5 and 8 neighbors on each side for each of the layers B5 and B4. The neighbor information at B1, B2 and B3 is adequate for 10 GeV/c or stiffer tracks. Thus the total number of inputs to this ASIC is 74 and the number of interconnections which must be made across all boundaries for track segment linking is 36.

The segment linker ASIC-2 will find the stiffest track candidate within its region of coverage and will create a 4-bit code for the p_t of the track. This code will be compared with the code from the adjacent ASIC-2 and the greater of the two will be passed along. After 6 clock cycles, the stiffest track found on each half of the readout board will be available at the 3rd and 6th ASIC-2 in the chain. These tracks are then reported to the trigger system for linking to the calorimeter and/or muon systems.

The combinatorial AND-OR logic of the segment linker will be designed to be remotely programmable within certain limits. Hence as luminosity increases and the demands on the trigger change, trigger thresholds can be changed dynamically to increase the effective p_t threshold.

Level 1 and Level 2 data storage

An additional task for ASIC-1 is to implement the Level 1 and Level 2 pipelined data storage functions and the interface to the DAQ system. Level 1 storage is implemented by a FIFO-like storage that runs at the 60 Mhz clock rate. This storage is constructed of 2 banks of RAM, each bank capable of storing data for the Level 1 Trigger delay time. For a delay time of 4 μ s, for example, each bank would need 256 storage locations. The two banks have a common address register, which is an eight-bit counter. Since both the read and write access to the banks is strictly sequential, the counter is simply incremented each beam clock cycle. While bank 1 of the RAM is being written with a set of data every beam crossing bank 2 is being read. When the address register turns over to zero, the banks are switched and the process continues, with

bank 1 being read and bank 2 being written. In this way, each bank sees access times no smaller than 16 ns, while providing simultaneous read and write capability.

A Level 1 Accept signal from the SDC trigger processor will cause the 64-bit word being read from the FIFO to be gated into the Level 2 storage FIFO. In the absence of a Level 1 Accept signal the word being read is discarded.

The Level 2 storage is set up as a true FIFO—a word written into the FIFO immediately falls to the lowest empty word location. Since the Level 1 Accept rate is relatively low compared to the beam clock, the Level 2 storage operates at a significantly lower speed than the Level 1 storage. A Level 2 Reject causes a word to be read and discarded from the Level 2 storage. A Level 2 Accept causes a word to be read from the Level 2 storage and transferred to the DAQ buffer FIFO. The depth of the Level 2 FIFO is expected to be 64 bits wide by 16 words deep. During final design, this may need to be adjusted when design decisions on the maximum latency of the Level 2 Trigger are finalized.

The DAQ buffer provides a holding register for data to be sent to the Data Acquisition System. This buffer is a small FIFO—64 bits wide by a depth of the order of two words. During final design, the depth will be adjusted as necessary. All data presented to the DAQ FIFO are eventually transmitted to the DAQ. The 64-bit word for the event currently being transferred is read from the DAQ FIFO and transmitted bit-serially to the DAQ interface on the Trigger and Readout board. Serial transmission is used to minimize the number of pins on the ASIC. These data are then transferred via the crate backplane to the DAQ system interface board.

VLPC support

Each Trigger and Readout board provides support to the VLPC cassette with which it communicates. The support includes temperature monitoring in the cryostat, pressure monitoring in the cassette, up to 96 individual bias voltages for the VLPC arrays, and power for the Amplifier/Shaper/Discriminators (ASDs). These services are provided via a single 128-conductor ribbon cable which connects one readout board to its corresponding cassette.

If needed, provision will be made to maintain bias voltages for the VLPCs by means of a capacitor array that will be refreshed by a RAM look-up-table and a Digital to Analog converter. The current needed for these bias voltages is extremely small. The refresh method is designed with a minimum of components, yet supports these small currents adequately. The look-up-table will be “down loaded” via the Control/Monitoring Network (Section 8.9), provided to all SDC subsystems.

Diagnostics

Via the Control/Monitoring Network, an artificial set of data may be loaded into the registers of the Trigger and Readout board. This provides a method of testing both the event data storage pipelines and DAQ interface, and also the segment finding and track linking capabilities. Because of the all-digital nature of the design of the board, the diagnostic data may be “played back” at full speed or may be single-stepped. This feature will be especially useful during commissioning and for diagnostics during normal running.

8.6.5. Interfaces to trigger and DAQ

Each of the 320 Trigger and Readout boards has a fiber-optic transmitter for sending trigger information to the SDC trigger system. Every beam crossing (16 ns), each card reports the 2 stiffest tracks found by the segment linkers on each half board. For each track, a 4-bit p_t value and 4 bits of Phi information are transmitted for each of these tracks. A p_t code of zero is transmitted if no track is found which satisfies the minimum p_t requirement.

Each crate also receives a set of signals from the SDC trigger system via the Level 1 Clock Control Board. These signals include beam clock referenced to the beam crossings, Level 1 Accept, Level 2 Accept or Reject, Reset, etc. These signals are distributed to the Trigger and Readout boards via the crate backplane.

Following a Level 2 Accept, each Trigger and Readout board sends its data to the DAQ interface card. A complete set of hit data (one bit per fiber) from a full crate (16 Trigger and Readout boards, each servicing 1024 fibers) occupies 16,384 bits—equivalent to 512 32-bit words. At a backplane transfer rate of one 32-bit word each 100 ns, 51.2 μ s are required to transfer an event from the Trigger and Readout boards to the CPU board (Section 8.9). At this speed, the crate can support a maximum Level 2 Trigger rate of 19,500 per second.

The data could be compacted by the CPU prior to transmission to the main data acquisition system. The decision to compact or not will be made on the basis of how the data might be used by the Level 3 processors. The bandwidth of the fiber link is great enough to allow uncompact data to be sent.

The CPU board also provides the interface for slow control functions. These include the monitoring and setting parameters such as VLPC bias voltages, temperatures, etc. and the downloading of pattern information to the segment-linking ASIC-2s.

8.6.6. Readout of fibers not involved with triggering

The U and V layers for the Scintillating Fiber Tracker do not contribute data to the Level 1 trigger processor. However, these fibers must still be read out as part of any event. This is accomplished using the same Trigger and Readout boards that are used for the other layers that do contribute to triggering. The only difference is that the segment-linking ASIC-2s are not installed on this board and no trigger system fiber-optic link is provided. The hit data from these fibers are then treated identically with data from trigger fibers as far as Level 1, Level 2, and DAQ buffering are concerned.

8.7. Intermediate tracking detector electronics

The Intermediate Tracking Detector (ITD) uses gas microstrip (μ -strip) technology. The design of the front-end electronics for the ITD is determined by the signal characteristics of gas μ -strips and by the trigger and data taking requirements. The high channel count results in the majority of the electronics issues being comparable to those of silicon detectors. The global design is very similar and is expected to use the same designs for many elements. In particular, the silicon data-driven pipeline chip, fiber-optic data transmission and external electronics needs very little alteration. An attractive proposal is to combine the external DAQ and Level 2 trigger systems for the silicon and the gas μ -strip system.

The R&D program to develop gas μ -strip electronics focuses on the areas that do not overlap with silicon. In particular, the pre-amplifiers require careful consideration. The packaging for the on-detector elements may be different since the electronics density is lower and allowable power consumption higher, which can affect the cost of the overall design as compared with silicon. Further, the ITD has a Level 1 trigger system, which is absent in the silicon detector.

In this section, the focus is on aspects of the ITD electronics which are different from the silicon; detailed discussion of pipelines and DAQ elements in common with the silicon tracker are deliberately omitted.

8.7.1. Detector and signal

Gas μ -strips are usually operated at fairly low gas gains. Here we assume a gas gain of 5×10^3 , which is typical of existing devices, and conservative in this context. The track length for a particle traversing the gas μ -strip varies between a maximum of 3 mm, when the track is perpendicular to the gas μ -strip tile, to a minimum of 200 microns when the track is parallel. High momentum tracks (> 5 GeV/c) will deposit ionization in the full 3 mm of gas. Because this is a relatively small path length compared to conventional proportional chambers and because we intend to operate these gas μ -strips at atmospheric pressure, a gas with good ionizing properties is needed. We assume the use of a gas similar in ionization properties to Xenon or DME that produces 200 ion pairs per cm. In that case, the signal after gas gain will be $3 \times 10^5 e$.

For a standard drift chamber gas with a drift velocity of 50 μ m/ns, electron collection times of up to 60 ns are expected. These numbers are conservative in that gas mixtures are available that can double

the drift velocity. The drift velocity determines the amount of charge that can be collected within a given time and hence affects the pulse height. In order to tag the beam crossing at the trigger level, charge is only collected for around 16 ns, which reduces the charge to approximately $8 \times 10^4 e$. Amplifier risetimes of 20 ns or better are needed.

To achieve good detection efficiency, the charge collected must be reduced by a factor of 0.55 to account for Landau fluctuations. A further factor of 0.5 must be included to allow for tracks that pass close to the center of a cathode and divide the resulting ionization between two adjacent anodes. The worst case charge to be detected is then Q_o ,

$$Q_o = 8 \times 10^4 \times 0.55 \times 0.5 = 2.2 \times 10^4 e.$$

However, this corresponds to only four primary ions. The clustering effect in primary production and fluctuations in longitudinal production location make it necessary to have a reasonable efficiency for single electron detection in order to avoid time jitter.

As stated above, the trigger and data taking requirements require that the timing of the output signal be resolved to within 16 ns. There are three main sources of timing jitter that need to be considered—the electron drift delay from track location, the drift delay from cluster production location and amplitude fluctuations.

- *Track Location Delay* is due to the time difference between pulses from anode-crossing tracks and cathode-crossing tracks. Using a drift velocity of $50 \mu\text{m}/\text{ns}$ and combining this with the results of the electric field calculations which have been expanded to include the calculation of isochrones, one obtains a figure of 3.5 ns delay for the 200 micron tiles and 7 ns for the 400 micron tiles. Thus there is a jitter of 3.5 ns on the inner radii tiles and 7 ns on the outer radii.

- *Cluster Production Delay* is due to statistical fluctuations in the position along the track at which the primary ionization is produced. Averaged over the Landau distribution of pulse-heights with the assumption of a 16 ns collection time, this effect adds $< 3 \text{ ns}$ jitter to the arrival time of the first electron at the anode.

- *Amplitude fluctuations* also cause jitter. Small pulses will cross the discrimination threshold later than large ones. Electronic noise also contributes to jitter. This contribution is minimized by using fast rise time amplifiers; ideally the rise time is determined by the collection time. If the range of charge collected in 16 ns is Q_o to $10 \times Q_o$ and the amplifier does not contribute significantly to the rise time, then in the approximation of a linear leading edge, the jitter from amplitude fluctuations is simply

$$[16\text{ns} \times Q_1/Q_o] \times [1 - 1/10] = 15\text{ns} \times [Q_1/Q_o]$$

where Q_1 is the discrimination threshold. A threshold at a level equivalent to $5000e$ threshold ensures that the jitter from amplitude fluctuations is small in comparison with other sources.

To minimize the tails caused by cluster formation location would require a threshold suited to detection of single electron tracks, which could only be achieved if very low noise front-end amplifiers can be developed. This threshold would be equivalent to $2500e$ at the preamplifier input. These considerations dictate that the amplifier noise must be no higher than $1200e$.

Detector capacitance has an important effect on the design of the front-end amplifier. Gas μ -strip capacitance is similar to silicon microstrips: $0.3 \text{ pF}/\text{cm}$ — 6 pF for typical strips.

In summary, the basic requirements for the amplifier and shaper are to achieve a noise of $< 1200e$ —preferably $800e$ —with a rise time of about 20 ns when the source capacitance is 6 pF . The power budget for this is $8 \text{ mW}/\text{channel}$.

8.7.2. Overview of on-detector electronics

The generic electronics chain for the gas μ -strip is shown in Fig. 8-41. The strips are partitioned into sets of 32; each set is fed into custom integrated circuit ASIC-01. This chip contains input protection, amplifiers and discriminators. The discriminators generate digital pulses representing hits. The hit information is passed to ASIC-02, which contains the Level 1 pipeline, read-out buffers and the Level 1 trigger-data reduction logic.

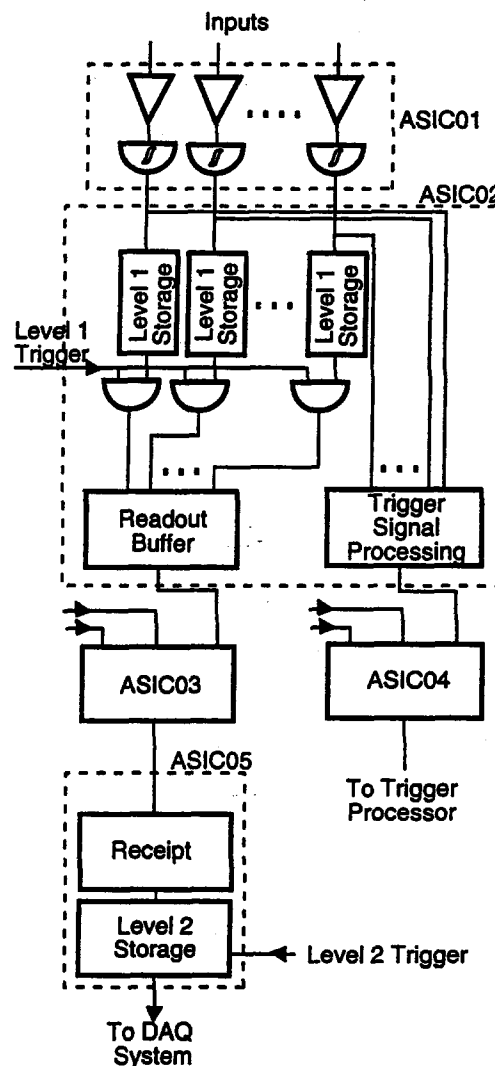


FIG. 8-41. The "generic" front-end electronics chain for the gas microstrip detector. ASIC's -01, -02, -03 and -04 are mounted on the detector. ASIC-03 transmits data on an optical fiber to ASIC-05 which is off-detector. ASIC-04 transmits trigger data on two optical fibers to the track finding logic located off-detector.

The pipeline is a content addressable memory (CAM) so that zero suppression is automatic. Data pass serially from ASIC-02 to local data collection chip ASIC-03, which multiplexes the data onto a fiber-optic link. The link moves the data off the detector and into receiver/buffer chip ASIC-05 where they are reorganized into a parallel format for storage in the Level 2 buffers and manipulation by the Level 2 trigger logic. In response to a Level 2 trigger the data are moved asynchronously into the output buffers, from which they are moved into the DAQ system. The trigger information by-passes the Level 1 pipeline and is

brought through the data reduction logic and multiplexers to Trigger Collection chip ASIC-04, where it is sent off to the main trigger electronics.

8.7.3. Front end layout

The substrate plate supporting the microstrips is extended beyond the gas volume by 2.5 cm and the front-end amplifier and pipeline chips are surface mounted onto this region. Data collection chips and fiber optic drivers are also connected in this region, but this electronics has low connectivity to the plate and is mounted on a separate flex circuit to simplify assembly.

The anode strips are terminated in bond pads $70\ \mu\text{m}$ by $100\ \mu\text{m}$ immediately after they exit the active area of the detector. The anodes are operated at ground potential and can be connected directly to the front-end amplifiers. The cathodes continue past the anode bond pads and join in groups of 32 which match the front-end chip segmentation. This is the analog "ground" and is attached to the analog ground in ASIC-01 via one (or more) surface-mount capacitors. The digital ground plane also extends under ASIC-01; this is continuous ground, common to all ASICs on a single detector plate. Plates are not connected to their neighbors except via the power line trees.

The outputs from ASIC-01 are bonded to one end of $70\ \mu\text{m}$ by $200\ \mu\text{m}$ pads and the inputs to ASIC-02 are bonded to the other end. The ground reference for this connection is separate to both the analog and the digital grounds. All bonding is done using standard wire-bonding techniques; the number of bonds is acceptable and adequate reliability should be achieved. If plastic substrates are used for the gas μ -strip, it will be necessary that it provide a reliable method of connecting to the rigid board containing ASIC-01.

The ASIC-02 outputs are connected to the flex read-out harness, on which the data collection chip (ASIC-03) and fiber-optic drivers are mounted using surface mount techniques. The flex is mounted flat on the detector plate so that the height of the flex components matches those surface-mounted directly on the plate.

Test pulse injection and high voltage supply are accomplished by extending the substrate plate beyond the gas volume in the opposite direction to the read-out electronics and using similar surface-mount techniques to connect the test-pulsing system and high voltage. The high-voltage ground is DC isolated from the analog and digital grounds.

8.7.4. Front-end amplifiers

It is essential to locate the front-end electronics as close to the gas μ -strips as possible where available space is severely restricted. This imposes a constraint on power consumption. It also results in a required packing density that necessitates the circuit implementation in custom integrated circuits.

Signal sizes and shapes are discussed in Section 8.7.1. They are similar to those of MWPC's. The relevant numbers are a rise time of 20 ns and an equivalent noise charge of $800e$ (1200 worst case) with a source capacitance of 6 pF.

An amplifier with approximately these specifications has already been designed by Bellazzini [50]. It is an 8-channel chip but has a noise equivalent of $700e$, a peaking time of 25 ns, a gain of 50 mV/fC and a power consumption of 5 mW (including the comparator).

Since the mean occupancy is low, the discriminator may be designed either to operate in a one-shot mode or in a time-over-threshold mode. This choice will be dictated primarily by the details of ASIC-02.

8.7.5. D.C. voltage supplies

Large amounts of low-voltage current must be supplied to the front-end electronics. The design goal is to keep the total power dissipation to less than 20 kW per end, but this still means that several thousand amperes of low-voltage current must be routed into the detector.

The design of the power busing must minimize material and avoid putting stress on the detector as a result of magnetic forces on large currents in power conductors. It is essential that grounding is done carefully. Ground loops are unacceptable.

The power distribution design is at a preliminary stage. Each detector plate is a leaf connected to a radial branch which extends from the outer radius of a power ring running round the outside edge of a superlayer. The rings are broken to avoid forming loops. The rings are not connected in z ; each superlayer is an isolated system. The advantage of this style of tree is that the heavy duty power buses (rings and z paths) are removed from the tracking volume.

Supply and return conductors are fabricated as a mechanically solid sandwich of aluminium and insulator. Since the supply and return currents are the same, there are no net magnetic forces on the power buses. Individual spokes will carry a total of around 25 A (supplying ten plates), so that they can be fabricated with a section of two 200 μm by 1 cm foils (on average) with a maximum voltage drop along a spoke of 0.1 V. Cooling is via conduction to the electronics cooling rings (Section 8.7.6). Since the current in the bus decreases linearly towards the inner radius, the spokes are fabricated as wedges, which yields an average of 0.05% of a radiation length per plane in the tracking volume.

8.7.6. Cooling

The electronics lie on rings, approximately 2.5 cm wide, with one ring per radial segmentation of the μ -strips. The bulk of the heating in the detector will be concentrated in these rings, each of which will dissipate 200 to 400 W (approximately 20 plates). Cooling rings are therefore required as an integral part of the mechanical support structure and are mechanically connected to the electronics. The design of these rings is not yet settled; however the power density in the ITD is much lower than in the silicon and the stability requirements are more relaxed. Thus, it is anticipated that the heat extraction problem is simpler.

8.7.7. Data transmission

For calculating data transmission rates, the mean occupancy of the strips at a luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ is assumed to be 0.4%. This figure assumes one-shot operation of the front-end discriminators. For a Level 1 trigger rate of 10^5 , the hit rate is 0.4 Mhit/s per tile (~ 1000 strips per tile). If zero suppression is used and if 10 bits are sufficient to encode the address of a hit, the data transfer rate will be 4 Mbit/s per tile. For a 32 channel read-out chip the data transfer rate will be 0.13 Mbit/s per chip.

To transfer the data over a few high-speed links requires multiplexing of data between gas μ -strip tiles. This has the potential of destroying the ideal grounding configuration. Alternatively, a large number of low bandwidth (*e.g.*, 60 MHz) fibers could be used. This has attractive advantages as it preserves the low-voltage power distribution tree structure and simplifies the multiplexing electronics. For further simplification, one fiber per 32 channel front-end chip is used; approximately 3×10^4 fibers would be needed in total. Without zero suppression, 32 bits of information are read out in response to a Level 1 trigger.

LED light sources are viable since the fiber lengths are of order 100 m. LED driven links, using 5 mA at 1 V, have been operated in a manner compatible with non-return to zero, strobed 60 Mbit/s links. For zero-suppressed data transmission, the mean power consumption per LED of 10.0 μW . The total power in the system is then nominally 10 W, which is manageable.

Fibers are grouped into ribbons which run along flex circuits attached to the cooling rings. The LEDs are an integral part of this construction and the only connection to the read-out chips is a single wire bond. The receiving diodes are also mounted as an integral part of the ribbon, in a similar manner to the LEDs; only electrical connections are made to the ribbon.

8.7.8. Trigger primitives

Because of the projective layout of the anodes in the ITD tiles, a simple algorithm requiring that three hits lie on a line in z - ϕ space is used to form track vectors for the Level 1 and Level 2 triggers. The major problem facing the Level 1 trigger for the ITD is to collect the very large amount of data at some central point in a sufficiently short time for them to be useful at the first level. The exact form of the ITD Level 1 trigger system is dependent on the bandwidth of the data transmission system that can be employed. Our current assessment is that some form of data compression is required on the detector before transmission to the central system which forms the Level 1 trigger. Even with this compression, tens of thousands of 60 Mbit/s fiber links are still required; however, our work on the development of low-cost low-rate fibers leads us to believe that this is a practical option.

The segmentation for the trigger is four bins in η in each of the two ITDs. Each of these η bins has a number of cells in ϕ which varies from about 1440 to 2880, depending on the radius, making a total of approximately 10800 cells per end cap. The cell size in the system ranges from a maximum of 6.3 mm to a minimum of 1.3 mm. The coincidence between the 3 superlayers would be implemented on 32 trigger cards (16 per crate) for each ITD. The crates are located at ground level in the counting house. Each crate is an SDC standard DAQ crate containing a DAQ CPU and a Level 1 Clock Control Board.

For presentation to the Level 1 trigger system, a single bit is set if there is a track above a p_t threshold in each of 64 $\Delta\phi = 0.1$ bins by 4 $\Delta\eta = 0.25$ bins for each end.

Data transmission

To form a useful Level 1 trigger, the data from approximately one million strips must be processed. This data must be made available to the trigger processor every beam crossing. If the mean occupancy is 0.4%, as discussed in the previous section, the data transmission rate is

$$60 \text{ MHz} \times 10^6 \text{ channels} \times 0.004 \times 10 \text{ bits/channel} = 2,500 \text{ Gbit/s.}$$

Associating two 60 Mbit/s fibers with each 32-channel chip involved in the generation of Level 1 trigger data gives 32,000 fibers, which is not impractical. Chips located on the inner ring of the detector may need to have three fibers. By associating fibers with chips, all need for multiplexing would be removed.

The primary problem with simply transmitting the data off the detector for manipulation in the trigger processor is that it is unclear what the effect of correlations in the data will be. Sixteen consecutive hits, requiring 160 bits, would take almost 3 μs to transmit, which is incompatible with Level 1 trigger processing. Until the correlation problem is resolved it is necessary to develop schemes for reducing the data before transmission. Two such schemes are described in the following section on triggering. One of these has been designed in detail and laid out in silicon. Fabrication of test chips is currently taking place.

Reduction schemes

Having extracted the data from the front end, the algorithm required to find z - ϕ tracks for the trigger is relatively simple and easy to implement. The proposed hardware works in one of three modes selected under software control.

Simple OR primitive

If the occupancy of the strips exceeds 0.5%, reduction mode one is used. A simple OR of the 16 strips is formed, the granularity of the detector is reduced accordingly and each fiber transmits hit-bits synchronously with the beam crossings. The mean hit rate per fiber will exceed 8%. The utility of this mode of operation has not been evaluated yet.

MPO primitive

For occupancies below 0.5%, reduction mode two, Matched Precision OR (MPO), can be used. For each bunch crossing, the data for 16 adjacent strips are examined. The data are then sent out as a serial string of bits which represent a one to one map of the hit pattern (16 bits), or as a shorter string of bits with fixed reduced precision (8 bit pattern of strip pairs, 4 bit pattern of strip quads, *etc.*).

Since the data are shifted out at 60 Mbit/s, many beam crossing occur before the trigger data for one crossing have been shifted out. It is possible to have a second beam crossing with hit data before the previous one is shifted out. To allow for this a four-deep buffer is used. A fifth hit crossing would simply be lost and represent a local trigger inefficiency. An additional 6 bits of time delay code are used to record the delay between event beam crossing and beginning of the data transmission. Adding start and stop bits results in 9 to 24 bits, depending on the level of ORing.

Acceptable momentum resolution can be achieved in the MPO scheme. A fixed number of strips is ORed before data transmission, but the number in the OR is tuned by location to match the trigger granularity required.

We make use of the simple relationship

$$\Delta\phi = \frac{0.3\Delta z}{z} \frac{r}{p_t}.$$

In order to have a constant resolution in transverse momentum, it is necessary to have increased accuracy closer to the beam line. Thus the natural way to divide up the planes of the ITD is to OR a small number of strips together to form a cell at the innermost radius, and a larger number at the outer radius. Suitable numbers would be one strip in the innermost radius and up to four in the outermost radius.

Having collected all the data from the ITD superlayers, the subsequent processing is relatively simple. The first step is to combine hit information for the first and last layers of the four-layer package. These are both radial layers and assumed to be separated by 2 cm. Each cell in the front layer would be combined with the appropriate cell in the back layer with a simple piece of circuitry which would demand that a hit be formed on a strip less than a certain distance from the straight line projection of the first hit strip. Demanding that this distance be less than 0.5 mm, for instance, would only be satisfied by tracks with $p_t > 1$ GeV/c at the innermost radius and 9 GeV/c at the outermost radius.

To define a track finding road width, assume that $\Delta\phi_1$ and $\Delta\phi_2$ are the maximum allowed deviation in ϕ hits between the first and third layers and between the first and second layers. Simple geometry indicates that the acceptance of the above algorithm as a function of $d\phi/dz = 1/p_t$ is 1 until the point $(\Delta\phi_1 - \Delta\phi_2)/\Delta z$, and then falls linearly to 0 at the point $\Delta\phi_1/\Delta z$. The linear range in p_t is controlled directly by the size of $\Delta\phi_2$. If for instance $\Delta\phi_2/\Delta z$ is chosen to be $(1/30)$ (GeV/c)⁻¹ and $\Delta\phi_1 = 2\Delta\phi_2$, 100% efficiency for tracks with $p_t > 30$ GeV/c is obtained and 0% for $p_t < 15$ GeV/c. Similarly $\Delta\phi_1 = 3\Delta\phi_2$ implies 100% efficiency at 15 GeV/c, 0 at 10 GeV/c, *etc.* Thus this level of processing makes simple coincidences between the three superlayers in order to give a simple trigger with a variable threshold.

VPT primitive

For occupancies below 0.5%, reduction mode three is an alternative. Variable Precision Trigger data (VPT) maximizes the information available for triggering by transmitting hit-density dependent precision data. The readout depends on the occupancy of the 16 adjacent strips:

- *One Hit Case.* If there is only one hit in the crossing, the hit location is encoded in 6 bits (lowest 4 bits are significant). This information is shifted out serially a fixed time after the beam crossing to which the data belongs.
- *Two Hit Case.* If two strips are hit in the crossing, the location of each is encoded with 3-bit accuracy (reduced precision with two strip resolution).

- *Three or more Hit Case.* For three or more hits, 2-bit accuracy is achieved by transmitting a simple one- to four-bit map of the hit pattern (reduced precision with four strip resolution).

An additional two bits are added to the serial data to indicate which of the three possible meanings is to be ascribed to the 6 bits of address data (the encoded hit strip pattern).

As in the MPO case, is possible to have a a second beam crossing with hit data before the previous one is shifted out. The same four-deep buffering, 6-bit time delay coding, and start-stop bits are used. The data transmission is 16-bit words.

For triggering the data are manipulated as follows. Neglecting the neighbor problem for clarity, the data corresponding to the three fibers that form allowable roads are brought together. The first and last plane information is directed into three synchronous steps of three parallel paths of look-up tables.

In the case of full precision data in both layers, only one of the three paths is used and only one of the three steps is needed. Full precision data consists of two four bit words and are used to address a RAM which contains all allowable combinations of two hits in the region and sets a bit in a 16 bit pattern which represents a one to one map of the strips in the middle layer. This bit must correspond to a hit in layer two for a good track to be found.

If layer one has one full precision hit and layer three has three reduced precision hits, then the data is directed into all three paths. Again, only one step is required in this case. Each of the hits in layer three is compared with the hit in layer one. All of the comparisons set bits in the same 16 bit pattern map representing the middle layer. In the case of three reduced precision hits in each of layer one and three, all three paths are used and each hit in layer one is compared with all of the layer three hits by using all three steps. The RAMs have ten inputs in total, 8 inputs for hit location and two to define the precision of the data.

The final step is to bitwise AND the data from the middle layer with the predictions made by the comparison of the first and last layers. A coincidence represents a found track and the precision code is a measure of track quality.

An integrated circuit has been designed and submitted for fabrication which performs part of the on-detector logic. The current version is single hit only, but will serve to verify that such logic can be incorporated on the front-end chips at an acceptable cost in power and silicon area.

Data receipt

At the receiving end of the trigger fibers, the data are first shifted in serially. For reduction modes two and three, the six address bits and the 2-bit interpretation code are then moved to an 8-bit wide 64-bit deep, synchronously clocked, shift register at the location in the shift register that corresponds to the time delay code. Thus if the time code was zero (no delay) the data would enter the shift register at the beginning; if the code was 64, it would be loaded at the last location.

At the end of the shift registers, all addresses belonging to a particular crossing are synchronized at a fixed time after the crossing.

Level 2 triggering

The Level 2 trigger makes use of all the ITD μ -strip information. The p_t of the highest p_t track is encoded into 4 bits for each of 1024 ϕ bins and 4 η bins for each end. The data collection chip sends the hit strip data to 10 off-detector crates, similar to the silicon crates, where it is buffered for DAQ read out and is also processed to form the Level 2 trigger data. These crates are located in ground level counting room. Each crate is an SDC standard DAQ crate containing a DAQ CPU and a Level 1 Clock Control Board. The beam crossing clock and Level 1 trigger signals are distributed to the on-detector electronics from these crates.

8.7.9. Pipelines

Independent of the shaping time for the front-end amplifiers, the pipelines for implementing the Level 1 storage and the ensuing logic will work with single-crossing resolution. The effect of shaping time is that the hit rate will be increased by a factor of three if time-over-threshold discrimination is selected.

The architectures suitable for the gas μ -strip ITD are identical to those useful for silicon, the only modification required for gas μ -strip use is that the pitch of the electronics be increased from 50 μm to 180 μm .

The architecture selected is a data-driven system based on a content addressable memory (CAM) for which test structures have been fabricated and test chips with a substantial fraction of the required functionality designed and submitted for fabrication. The test chip has 64 channels, which is already more complex than that required here. The chip has a version of the VPT data reduction scheme for use in Level 1 trigger primitive formation, as described above.

8.7.10. Status

Work pertaining to the gas μ -strip amplifier is going on in several groups worldwide. Several of these studies will have results available towards the middle of 1992. Designs specific to the ITD will be started once the generic R&D that is ongoing indicates the best direction to proceed.

The prototype Level 1 pipeline chip (with VPT data reduction) will be available in April 1992. After evaluation of this, we anticipate reworking the VPT to include the MPO scheme and producing a version of the chip with density suited to both silicon and gas μ -strip systems.

Low-speed fiber optics seem to be radiation hard and favorable in terms of cost. Trigger system design will be done based on the assumption that no pitfall is discovered in the ongoing program to develop this system. A version of the fiber optics suited to the ITD will be developed in the first half of 1992.

8.8. Silicon tracker off-detector electronics

The major portion of the front-end electronics for the silicon tracking system is located on-detector, bonded directly to the silicon strips. Following each accepted Level 1 trigger, the on-detector electronics transmits its data via optical fibers to a second portion called "silicon tracker off-detector electronics." This section describes the off-detector electronics, along with the transmission of data from the readout chips to the off-detector electronics. The on-detector electronics is described in Chapter 4.

Data transmission

When transmitting data, the front-end chips, which each service 128 silicon strips, use one of two formats. If the number of hits within the group of 128 is less than 10, hit addresses are encoded and transmitted serially. In the case of more than 10 hits, a string of 128 bits is sent, with the hit locations set to "1" within the string. This format ensures that the maximum time required by a front-end chip to transmit data for an event is 2 μs . Every chip sends a header word in response to a Level 1 trigger which starts with a marker, is followed by a format defining bit, and a 7-bit code that gives the number of data words in the string (which may be zero). The multiplexing scheme collects such strings and uses the same format in transmitting the data, adding a chip address to each word group and to each 128 bit group, but eliminating zero word headers. In the case of no data, a header is sent with a zero word count. When data are present, the header contains the number of active front-end chips, in exact analogy to the front-end chip format.

Two schemes for data transmission between the on-detector readout chips and the off-detector electronics are under development. One employs printed aluminium cables to create a high degree of local multiplexing and to route the data into high-speed laser links. The second uses a highly parallel approach in which a 60-MHz LED-driven fiber is dedicated to each front-end chip. Intermediate solutions involving fibers shared between a small number of front-end chips are also under evaluation.

Readout using high-speed fiber drivers

In this scheme, signals from a given detector segment are transmitted outside the detector by optical fibers operating at 300–500 Mbit/s after significant multiplexing. Since high-speed optical transceivers tend to be power-hungry due to the high-speed serialization circuitry and bulky compared to the strip readout ICs, they are situated at the outer radius of the Silicon Tracker volume. Signals from the individual detector segments are brought out radially to the optical transceivers by low-mass ribbon cables with aluminum conductor that carry both power and signals. Data on these lines are transmitted on a 12-bit wide bus at 30 or 60 Mbit/s (to be determined), allowing rapid readout of high-density jets concentrated in one module.

Since only a limited number of readout ICs can be daisy-chained with reasonable on-chip drivers, Bus Selector Chips are used to organize the readout of no more than 15 readout ICs. The segment bussing brings the signals to the end of the barrel, from which the cables are brought radially to the outer shell of the Silicon Tracker, where the optical transceivers are located. Bussing in the disks is organized in a similar manner.

Two possible locations of the optical transceivers are being considered. The first option locates the transceivers at the end of the Silicon Tracker volume where they are accessible for maintenance. The second option reduces the length of metal transmission lines by placing the optical transceivers at the end of the barrel and at the outer radius of the disks. In the first option the metal lines would consist of two parts. Ultra-low-mass cabling would be used to bring the signals from the detector modules to the outer radius, at which point the larger available area allows the use of wider conductors with less dispersion and a larger cable geometry (ribbon width and dielectric thickness scaled to preserve the characteristic impedance). A more detailed analysis of cable characteristics together with signal drivers and receivers is needed, but initially the first option seems feasible and is more attractive because of improved maintainability.

Readout using low-speed fiber drivers

LED-driven links based on existing low-cost components are being developed. Links have been operated in a NRZ mode at 63 Mbit/s. Radiation tests done by the manufacturers show adequate resistance to neutron fluences up to $10^{14}/\text{cm}^2$ and almost no degradation with photon doses of 1 Mrad. Both LEDs and PIN diodes are radiation hard. The system is low mass; at the detector end the only drive circuitry required is the diode itself and possibly one bipolar transistor. These are surface mounted on a flex circuit, which is an integral part of a fiber harness containing a group of up to 16 fibers. In tests to date, the fiber to diode alignment was accomplished by a low-mass photo-fabricated copper "block". Facilities are being upgraded to allow aluminium to be used, however even with copper, the total mass of the detector end system is a very small fraction of the silicon wafer thickness ($< 8\%$). Power consumption is negligible since the LEDs are off when not transmitting a "1"-bit and require of order 10 mW during transmission of a "1". The estimated cost per unit bandwidth is less or comparable to high-speed links, in part because the alignment problem is much simplified.

Associating one such fiber link with each front-end chip minimizes the bussing of high-speed signals, and hence simplifies the front-end DAQ architecture considerably.

Control and clock signal are also routed into the detector on fibers. All control information is routed as serial control words through the Level 1 trigger fiber, which otherwise is not in use during control operations. A continuous string of 10 "1"-bits on the Level 1 input is detected by the front-end chips and the following 8 bits are interpreted as a control word. No high-speed signals are routed on electrical buses, which contain only the power and DC analog control levels. These buses are, therefore, simplified in comparison with the high-speed driver case, but are still required.

In addition to one read-out fiber per chip, each detector wafer has associated with it four incoming fibers: two copies of the clock and two copies of the Level 1 trigger. These are translated to electrical signals and routed to the chips on the wafer via a four-line local bus. The redundancy in clock and Level 1 signals is incorporated to increase system reliability. In addition, if a readout fiber fails, the data from a chip can be routed through the adjacent chip into the adjacent fiber.

The number of fibers in this scheme can be reduced by *a priori* having front-end chips share read-out fibers. If the degree of multiplexing is low enough, the multiplexing functions can be incorporated into the front-end chips.

Data receipt

Figure 8-42 shows the overall structure of the off-detector electronics which receives the data from the fibers. The major tasks this electronics is to assemble, re-format and organize the incoming stream of data into data elements that are understandable by the Level 2 trigger processor and by the DAQ system. After the optical fiber receiver, the incoming data are in the form of parallel streams of asynchronous, serially encoded, data with a bit rate of 60 MHz. The number of such streams is determined by the details of the data transmission scheme employed in the final design and does not affect the basic architecture.




FIG. 8-42. Block diagram of the Silicon Strip Detector off-detector electronics.

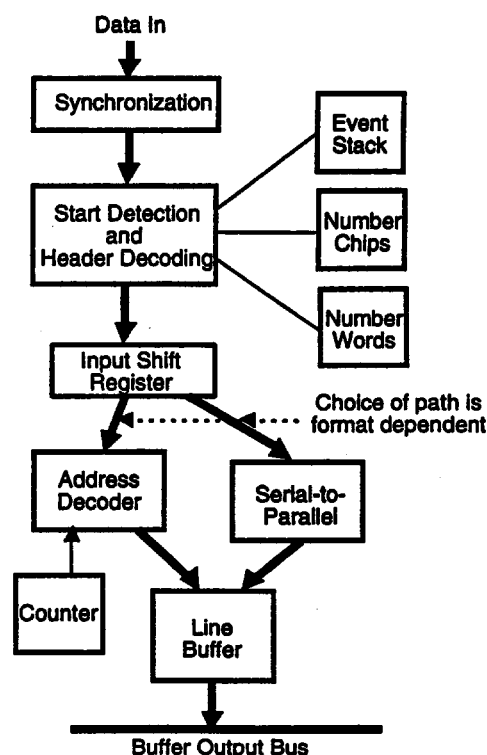


FIG. 8-43. Details of the Data Receipt block.

The first element in the chain is the Data Receipt block where the data are synchronized to the system clock (60 MHz), translated from serial to parallel format and stored in line receiver buffers. From here the information is grouped according to $\eta - \phi$ bins suited to the Level 2 trigger logic. The fibers themselves have sufficient segmentation that information does not need to be moved on the basis of the hit location, but exclusively on the basis of the data stream from which it originated. Once collated, the data are moved in parallel to both the Level 2 storage and to the trigger system, which is designed to accept an asynchronous data stream.

Figure 8-43 shows a block diagram of the Data Receipt logic for one input data stream. The data are first synchronized with the system clock. This is accomplished by routing the data stream through a variable number of buffer gates under the control of a phase-sensing circuit shown in Fig. 8-44. The length of the fibers and delays through any logic prior to this circuit are rendered irrelevant by this step since the




FIG. 8-44. Logic for automatic synchronization of data streams with system clock.

transmission is asynchronous. The only necessary condition is that the data stream and fiber logic is driven by the same clock as the rest of the DAQ system.

After synchronization, the data are directed into a shift register; the readout process is started when a non-zero bit is detected and the number of chip data sets to record is split off for control purposes. Each group of words is shifted-in serially and moved to the line buffer as a parallel word with the address expanded to include the chip address. Each 128-bit hit pattern is encoded as addresses and shifted to the buffer so that after the last chip group has been read in, all the hits corresponding to one event in the area serviced by the data stream are stored in the line buffer.

When all the data for an event have been read in, an entry is made in the event stack containing the number of data words and the starting location in the buffer. Events with no hits also result in an entry, but the starting location is a code indicating no data.

Line buffers share an output data bus inside the receiver chips and buffers corresponding to a trigger $\eta - \phi$ region share a data bus across chips. Data are moved to the Level 2 trigger and to the Level 2 storage simultaneously along this bus under local control. Data from all stacks must be shifted before any stack is repeated to ensure events are properly aligned. Since all Level 1 triggers are processed sequentially, if an event stack has any complete events recorded, the oldest one is the next to be processed. All stacks will contain data for each event since zero word headers are written.

Data remain in the Level 2 storage, stored as complete events by $\eta - \phi$ region until the Level 2 Reject or Accept is received. These arrive in order, so that the oldest event is either shunted to the output buffer or discarded. From the output buffer the data is moved synchronously to the processor farm.

Physically the hardware is an integral part of the silicon Level 2 trigger processor and housed in the same SDC standard crates.

8.8.1. Triggering

Trigger studies are emphasizing the reconstruction of high transverse momentum tracks for a second-level trigger. This requires readout of the silicon after a first-level trigger. A specific second-level trigger scheme using hit patterns has been defined and an estimate of the number of trigger cards completed. This scheme will be looked at carefully, including simulation studies, next year. At lower priority, schemes for generating a contribution to the first-level trigger are also under study.

The second-level trigger for the silicon detector allows the full granularity and momentum resolution of the detector to be utilized at an early stage in the trigger. It provides a powerful aid in rejecting photon conversions which form the main background to electron triggers at the first level. Using the full ϕ granularity, a momentum resolution of $\Delta p_t/p_t^2 \simeq 1 \text{ (TeV/c)}^{-1}$ is possible.

The Level 2 trigger is designed to accept an input rate of up to 10^5 Hz. Data from the silicon tracker in the form of a list of hit strip numbers is multiplexed up to the processors on optical fibers. The data are split, one copy being stored in a Level 2 data buffer for possible DAQ readout, while the other is routed to the Level 2 trigger logic. The data from a given first-level trigger is completely assembled within about 2 μ s of the first-level trigger being issued.

The second-level trigger for the silicon tracker is based on simple pattern recognition algorithms operating on tables of hit positions in a two-dimensional space. In the case of the barrel silicon modules the two dimensions are the radial distance of the silicon layer from the origin and the azimuthal angle ϕ . In order to utilize the full precision of the silicon detector while keeping the amount of hardware to a minimum we adopt a two-stage approach to the pattern recognition problem. We first attempt to recognize regions of interest in the table by means of a "coarse scan". Having established areas of the two-dimensional space likely to contain high p_t tracks, we then carry out a second, "fine scan", in which these areas are examined with the full granularity of the silicon detector. This allows a high precision measurement of the transverse momentum of those tracks found by the pattern recognition.

In the coarse scan the data from the barrel region is tabulated in 1920 bins in ϕ . This bin size is chosen to approximately match the ϕ granularity of the calorimeter and to keep the number of inputs to the processors to a manageable number while maintaining good resolution in transverse momentum. The two-dimensional table is formed by combining the 1920 bins from 8 layers of the silicon ordered by increasing radius. Clearly in such a table high transverse momentum tracks have a very small or zero variation in ϕ bin as a function of radius, whereas low p_t tracks bend considerably as they pass through the eight silicon layers. The trigger processors must be able to discriminate between the presence of a track and of random noise in the table and, in addition, to discriminate between high and low transverse momentum tracks. In the conceptually simple case of infinite p_t and full efficiency in each layer of silicon, this problem reduces to being able to recognize tables with eight hits in the same ϕ bin. Allowance for inefficiency and for finite p_t requires somewhat more complicated algorithms, but clearly the pattern recognition problem is at least conceptually rather simple.

An almost identical approach is used for the endcap modules of the silicon detector, except that, rather than filling a table of ϕ versus radius of the silicon layer, we construct a table of ϕ against z , the distance along the beam line. The same granularity in ϕ is used as for the barrel. Again, high transverse momentum tracks have a small variation in ϕ as z changes, whereas low p_t tracks curve considerably. The pattern recognition problem is therefore exactly analogous to that in the barrel once the z versus ϕ table has been constructed, so that identical trigger processors can be used. The case of tracks which cross the barrel-forward interface needs further study and is an important class of tracks.

In both the barrel and endcap modules there is a further subdivision of the data in bins of pseudorapidity η . These bins correspond in the barrel to the length of silicon wafers bonded together to form a readout unit which is expected to be two wafers. These 6 η subdivisions are somewhat coarser than that of the calorimeter, but should still allow satisfactory matching of tracking information and energy deposition.

Two possible hardware implementations of the silicon Level 2 trigger are discussed in detail in Ref. 72. The current estimate is that the triggering system can be built in approximately 10 crates. The crates are located in the ground level counting house. Each crate is an SDC standard DAQ crate containing a DAQ CPU and a Level 1 Clock Control Board. The beam crossing clock and Level 1 trigger signals are distributed to the on-detector electronics from these crates. It is likely that track segments found in the outer tracker by either Level 1 or Level 2 logic can be incorporated into the silicon Level 2 trigger in a manner analogous to the silicon hits. With additional information, the pattern recognition capabilities of this trigger may be strengthened and the sensitivity to the location of the beam spot will be reduced.

8.9. Data acquisition system

The major job of the DAQ System is to collect from the front-end and trigger systems the digital data that represent the events selected by the Level 1 and Level 2 Trigger Systems and to transfer this data into an array of processors (Online Processor Subsystem) for online filtering (Level 3 triggering). It is the function of the Online Computer System (Chapter 9) to transfer processed event data from the Online Processor Subsystem to permanent storage for later offline analysis. A complete list of the responsibilities of the DAQ System is given below:

- Provide software and hardware to receive information regarding each Level 2 Accept trigger; to collect event data fragments from front-end and trigger systems; to concatenate these fragments into complete events ("build" events); and to transfer the events into an online array of processors.
- Guarantee the integrity of the collected data.
- Provide an online array of processors to filter collected events (the Level 3 trigger procedure). Communicate with the Online Computer System to transfer events satisfying the Level 3 trigger criteria onto online permanent storage for later analysis.
- Provide software and networks for downloading, controlling and monitoring the operation of the front-end and trigger system components housed in readout crates, DAQ System components, and other remote (on-detector) front-end and trigger system components.
- Provide powered and cooled crates and support electronics to house front-end and trigger system electronics.
- Provide powered and cooled racks and rack protection software and hardware for both front-end and trigger system components housed in readout crates and DAQ System components housed either in readout crates or elsewhere.
- Provide software and a network interface to the Ancillary Controls System from the rack protection electronics.
- Provide runtime user interface software for such tasks as runtime control, control of the online array of processors, histogramming, event displays, calibration data bases and consumer processes.
- Provide control, monitoring and downloading software for supporting the use of memories and intelligent processors in the DAQ System, for supporting "slow" (non-event) data acquisition and for supporting the control and monitoring of adjustable devices.
- Provide software for quickly detecting system errors and recovering from these errors.
- Provide stand-alone (portable) DAQ systems for use in test beams and in lab tests of front-end electronics and detector systems.

The DAQ System interfaces to the front-end systems for event data collection, to the trigger systems for both event data collection and event data flow control, to the Online Computing System for event data storage as well as overall setup and runtime control and monitoring of the DAQ System, and to the Ancillary Controls System both for control and monitoring of the DAQ System's electronics protection circuitry and additional overall control and monitoring.

This section begins with discussions of the functional requirements of the DAQ System. This is followed by an overview of the proposed DAQ System architecture including a block diagram. Emphasis in this section is placed on the description of the architecture of the DAQ System, first at the block diagram level and later covering details of the various system components.

This section also contains a discussion of how the system can be partitioned for various modes of operation, descriptions of the components of the system, system software, event data flow control, error detection and recovery and a discussion of stand-alone (or portable) DAQ systems.

The section continues by listing quantities and locations of the components of the DAQ System. The installation, integration and testing of the entire DAQ System and some ongoing simulation activities are then discussed.

The section concludes by listing the system's development and commissioning schedules, milestones and critical path items.

8.9.1. Data acquisition system requirements

The Level 1 and Level 2 Trigger Systems determine which events are to be read out by the DAQ System. The DAQ System collects data from the selected events and delivers these data to its online array of processors (Online Processor Subsystem/Level 3 Trigger) where the final stage (Level 3) of event selection is performed. Upon receiving event data fragments from the front-end and trigger systems, the DAQ System assembles these data fragments into coherent events in the Online Processor Subsystem. After Level 3 Trigger criteria are applied to these events and upon request from the Online Computer System, the DAQ System transmits selected events from the Online Processor Subsystem to the Online Computer System's permanent online storage facility.

The requirements for the SDC DAQ System fall in four broad categories. First are the quantitative requirements such as those on the bandwidths of data paths and the number of data sources to be accommodated. Next are the important architectural requirements that allow front-end and trigger systems to be operated together or independently. Next are the requirements for initializing, controlling and monitoring the system. Lastly are the qualitative requirements on items such as reliability, tolerable error rates and maintainability. A more complete discussion of the requirements for the DAQ System is given in Ref. 51.

Figure 8-45 shows the basic data paths for the SDC detector and throughput requirements at various stages of the trigger and data acquisition systems. These systems are based on a hierarchical series of triggers, Level 1, 2 and 3. The generation of Level 1 and Level 2 triggers is described in Section 8.2. At each trigger level, the corresponding trigger is expected to achieve an event rejection within the range shown in the figure. For example, a Level 1 rejection range of 1 K to 10 K means that from one event in a thousand to one event in ten thousand is accepted. The event rates, bandwidths and deadtimes shown in the figure drive the requirements for the DAQ System. An aggregate of greater than 10^6 channels provide information at a rate that requires a overall DAQ System bandwidth of 1 Gbyte/s (typical), and 10 Gbyte/s (maximum).

Data flow and event collection

The fundamental requirement on the Data Acquisition System is that it be capable of collecting, from all front-end and trigger system data sources, all event data for those events selected by the Level 2 Trigger System, while imposing less than a specified amount of dead time.

The dimensions of this problem are outlined in Table 8-9. Here, information for each front-end and trigger system that generates event data is listed. The column labeled "No. Crates" gives the number of readout crates required to house the modules associated with each system. All of the crates are of the same design, the SDC DAQ standard readout crate design. Note that the Central Calorimeter Switched Capacitor Array and Shower Maximum detector systems share the same crates. Each crate transmits event data to the DAQ System over one or more optical fiber links. The total number of fibers for each system (or, in the case of Central Calorimetry, for both systems) is listed in the column, "No. Sources."

As shown in Table 8-9, the total number of crates that generate event data is 274. Some crates generate more than one event data stream ("source"). The total number of independent front-end and trigger system data sources is 394. Typically, each data source will prepare and transmit to the DAQ System, a packet of data for each event for which a Level 2 Accept trigger is generated. All these sources together are expected to generate approximately 1 Mbyte of event data per event. Since the maximum Level 2 Accept rate at design luminosity is expected to be 10^3 per second, the total available bandwidth for moving data from the 394 sources is required to be at least 1 Gbyte/s.

Other requirements are derived from the expected event size, the Level 1 and Level 2 trigger rates, and the deadtime requirement of less than 10% at design luminosity. A list of important requirements is given in Table 8-10.

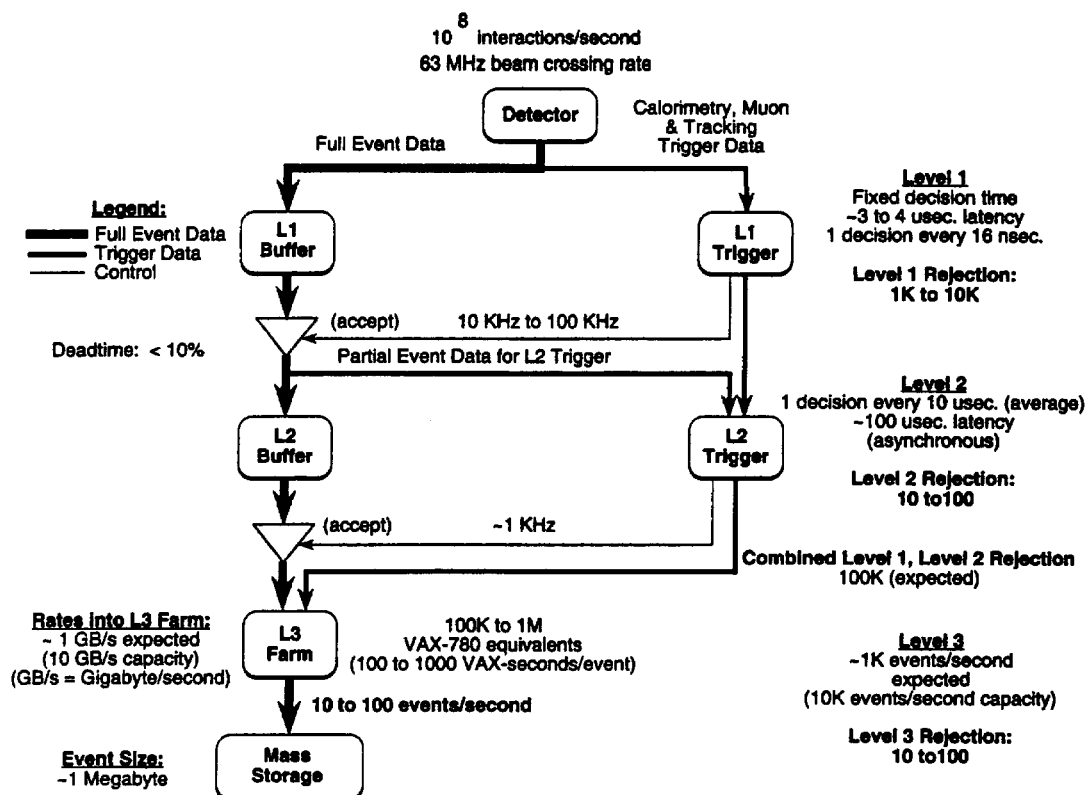


FIG. 8-45. SDC event flow diagram.

The major design parameters for the initial system may be summarized as follows: 394 event data sources; 1 Mbyte of data per event; 1 Gbyte/s (1000 events per second) through the Event Builder Subsystem and into the Online Processor Subsystem; and 10⁵ VAX-11/780 MIPS processing power in the Online Processor Subsystem. The system must scale in a modular way to accommodate an upgrade in luminosity, which will result in an event data rate of 10 Gbyte/s (10⁴ events per second) through the Event Builder Subsystem and into the Online Processor Subsystem, and 10⁶ VAX 11/780 MIPS processing power in the Online Processor Subsystem.

Data integrity

The DAQ System must insure the integrity of the event data stream. That is, the data must have been transported from the front-end and trigger systems to the Online Processor Subsystem without error in either data content or identification of origin.

System partitioning

Partitioning is the concept of logically dividing the electronics of the detector into two or more independent subsystems, each with independent control, triggering, deadtime and data acquisition facilities. When partitioned, each subsystem must be able to operate as an independent, stand-alone entity, free from interference from other subsystems. The capability to partition the system is an essential requirement, both during commissioning and during routine operation, for a number of important reasons. For example, it will be impractical, if not impossible, to commission the detector in the time allotted if the various front-end and trigger systems must wait their turn to use the complete Data Acquisition System to execute calibration, monitoring, diagnostic testing and data acquisition tasks. Partitioning will allow the individual systems to do these tasks simultaneously.

It will be possible, under computer control and without massive re-cabling, to assign individual crates or entire detector front-end systems to a partition. It will also be possible to simultaneously activate several tasks in a single partition to cause different data sets to be read out for different types of triggers.

Table 8-9

Number of crates equipped with DAQ CPU modules and number of event data sources for each front-end and trigger subsystem.

Front-End/Trigger System	No. Crates	No. Sources
Inner Tracking Silicon Strips and		
Inner Tracking Silicon Pixels	10	10
Intermediate Tracking Gas Microstrips	10	10
Outer Tracking Straw Tubes ¹	8	32
Central Calorimeter (SCAs) and		
Calorimeter Shower Max ²	96	192
Central Calorimeter (Flash ADCs) ³	96	96
Forward Calorimeter	2	2
Muon	64	64
Level 1 Trigger (≤ 1 kbyte/crate)	59	59
Level 2 Trigger (≤ 1 kbyte/crate)	25	25
Totals:	274	394

¹ Each of the 8 crates has four modules each of which contains a data source that transmits event data directly to the Event Builder Subsystem.

² Each of the 96 crates has one Data Collection module with one Calorimetry and one Shower Maximum detector source. Event data are transmitted directly to the Event Builder Subsystem from the Data Collection modules.

³ Central Calorimetry (Flash ADCs) are listed for completeness but, as was the case for SDC detector costing, are not counted in totals since the Flash ADC option, if chosen, would take the place of the SCA electronics.

Table 8-10

Data flow and event collection requirements.

• Maximum Level 2 Trigger System input rate:	100 kHz
• Maximum Level 3 Trigger (Online Processor Subsystem) input rate:	10 kHz
• Maximum bandwidth through Event Builder Subsystem	10 Gbyte/s
• Minimum processing power in the Online Processor Subsystem	10^5 VAX 11/780 MIPS
• Minimum archival recording rate	100 Mbyte/s
• Expected event size	1 Mbyte
• Maximum readout deadtime	10%
• Maximum deadtime due to DAQ System errors/downtime	5%

Data acquisition system control and monitoring

The DAQ System must also provide a set of services to control and utilize the calibration and monitoring facilities built into the various front-end and trigger systems, to handle error and alarm systems, to record various "non-event related" data streams (magnet currents, high voltages, accelerator parameters, etc.), to control the various processors and modules in the system and to deliver events to consumer processes for miscellaneous on-line monitoring purposes. This set of tasks is collectively performed by the DAQ Control/Monitoring Network (Subsystem) under supervision of the Online Computer System. A more complete list of the functions of the DAQ Control/Monitoring Network is given below:

- Provide the capability of setting up the entire detector into a known (and recorded) condition.

- Track and log the operation of the DAQ System as well as the various front-end and trigger systems.
- Allow for non-event data streams such as calibration data acquired by the front-end systems, database accesses, *etc.*
- Record the conditions under which the data were taken.

Detect and record all error conditions.

- Maintain a set of prioritized alarms to shift personnel regarding detector status.

The DAQ Control/Monitoring Network has to support two types of data traffic. One type of data traffic requires relatively small amounts of data (*e.g.*, voltages, temperatures, *etc.*) to be made available to the control and monitoring processes, while the other less frequent type of data traffic requires large data sets to be either downloaded or retrieved (*e.g.*, during system debugging). More specific requirements are given in Ref. 51.

Software environment

It is likely that the various processors in the DAQ System require several types of operating environments. The operating systems and network protocols are required to be commercial standards. For example, the processors charged with readout responsibilities may run under a real-time operating system environment such as VxWorks and the processors in the Online Processor Subsystem may run under an operating system such as UNIX. Processors connecting to the DAQ Control/Monitoring Network or the Ancillary (Slow) Controls Supervisory LAN may run under standard network protocols such as TCP/IP. The application programs written by various groups must conform to prescribed standards of coding and revision histories.

Commissioning the system

During the initial commissioning of the DAQ System (installation, integration and testing at the SDC detector site), collaborators responsible for various front-end and trigger systems will need to test their electronics, software and other detector components independent of each other's systems. Separate workstations and an independent working data acquisition and control/monitoring environment will be initially provided to these collaborators. This environment will be a subset of the final Data Acquisition System. Later these workstations will become part of the final system or possibly remain as they are initially configured for future needs.

System growth and scalability

Inevitably, the data acquisition requirements will change during the useful life of the SDC detector. For example, additional channels may be added to existing front-end systems or additional front-end systems may be added to service additional detectors. The operating parameters of the storage ring (frequency of beam crossings or beam intensity) may change and thus the required throughput of the data acquisition system will change. Careful attention has been taken to ensure that the DAQ System can gracefully "scale", in a modular way, to accommodate different numbers of front-end channels and the concomitant data throughput.

A key element in this scaling capability is the modular design of the Event Builder Subsystem which enables this subsystem to increase its throughput linearly (true parallel event building) with increases in the number of input sources.

Technological advances

Care has been taken to propose an architecture for the Data Acquisition system that will satisfy the experimental needs for the useful lifetime of the SDC detector. Advances in technology will inevitably occur and, over time, these advances will be installed in various parts of the hardware and software of the system. However, if the basic architecture design easily adapts to these changes, substantial cost savings will be realized in the future.

The point-to-point links used for digital data transmission in all parts of the DAQ System (and particularly in the switching network portion of the Event Builder Subsystem and the Online Processor Subsystem) are implemented with industry-standard protocols. It is expected that the improved data link components that result from advances in technology will be designed to utilize these same protocols. This will permit the performance of the system to be upgraded with time with minimal or no significant architectural modifications.

Reliability and maintainability

Care is being taken to ensure a high level of reliability in the DAQ System, both in terms of component reliability and in terms of error-free data transmission. The components in the DAQ System are required to be designed using standard industry practices for reliability. Proven industry standards and well-supported commercial software and hardware components have been proposed wherever possible. Redundant components or data paths are planned in certain critical locations. Diagnostic features such as error detection and correction schemes and diagnostic buffers at various locations throughout the system are being implemented. Since these features add complexity and cost, attention is paid to achieve the proper compromise between costs and benefits.

In order to minimize downtime, the DAQ System should not only be able to detect any occurrence of errors and the sources of the errors but also be able to correct the errors quickly once their sources are known. The DAQ System is being designed with an emphasis towards implementing the proper set of diagnostic features to achieve these goals.

The requirements on the SDC Data Acquisition System that address the reliability and maintainability concerns are:

- The data transmission links used for event data are completely separated from those used for transmission of system control and monitoring data.
- Data integrity checks (using, *e.g.*, CRC, or other types of check words, and possibly correction codes) are part of the data transfer protocols at all levels of the DAQ System. Checking of ID parameters, such as event number or beam crossing number, verifies the correct routing of event data fragments to their final destination.
- Facilities for logging error conditions and event data histories are implemented at several system levels (the readout crates; the Event Builder Subsystem and the Online Processor Subsystem).
- Facilities for error logging functions, including event data flow control histories, are implemented in the Event Data Flow Control Subsystem.
- Industry-standard (and proven) data transmission links and industry-standard busses are used throughout the system.
- All electronics enclosures are provided with temperature and humidity control.
- System hardware is implemented with modular components wherever possible to permit rapid correction of detected malfunctions.

8.9.2. Overview of the data acquisition system architecture

A block diagram showing the major components of the SDC Data Acquisition System and the immediate neighbor systems to which it interfaces is shown in Fig. 8-46. In this diagram, event data flow from left to right.

The interface between the front-end and trigger systems and the DAQ System is distributed throughout an ensemble of 274 readout crates (cf. Table 8-10, symbolically represented by the three sets of crates on the left of the figure. Housed within the readout crates are components of the various front-end and trigger systems that collect signals from detector elements, process them in response to the Level 1 and Level 2 triggers, and deliver digitized versions of these signals (event data) to the DAQ System.

Each readout crate contains two modules that are part of the DAQ System. The first module is a Data Acquisition System processor module (DAQ CPU) that event-by-event reads the event data fragments in the entire crate. It then re-formats as necessary and transmits these data (through the second module) to the Event Builder Subsystem over an optical fiber. (Some front-end systems have data collection buffer modules which transfer their event data fragments directly to the Event Builder Subsystem, bypassing the DAQ CPU module.) The DAQ CPU is also a node on the DAQ Control/Monitoring Network.

The second module is a Crate Adapter/Interface module used to physically adapt the smaller commercial DAQ CPU module into the larger readout crate, and to house logic used to transmit the event data fragments to the Event Builder Subsystem. Another use of this interface logic is to extend a subset of the DAQ Control/Monitoring Network to parts of front-end systems not housed in crates and located in hard-to-reach areas within the detector.

Each readout crate also contains a Trigger/Gating module which is part of the Trigger System. (It is referred to as the Level 1 Clock Control Board, L1CCB, in Section 8.2.) This module receives Level 2 Accept information over a high-speed fiber optic cable and passes this information to the crate's DAQ CPU, where it initiates the readout of data for the selected event. The DAQ CPU uses this module to inform the Trigger System of conditions such as an impending buffer overflow via another high-speed fiber optic cable. Such conditions cause the trigger system to momentarily inhibit the generation of triggers.

Continual streams of event data are transmitted from the 274 Readout crates to the Event Builder Subsystem 394 optical fibers each of 1 Gbit/s capacity. The streams consist of successions of event data fragments, each fragment containing the crate's set of data for an event selected by the Level 2 Trigger System, together with identifications of the origin of the data and the event number.

The Event Builder Subsystem receives the streams of event data fragments from all its 394 inputs and builds these fragments into complete arrays of data, each corresponding to a single event. This event building is accomplished with the aid of data balancing and input queueing logic and a commercial switching network within the Event Builder Subsystem. Simultaneous streams of totally built events are transmitted in parallel to the Online Processor Subsystem.

A Level 3 Trigger inspection of the data is executed in the processor array of the Online Processor Subsystem. Event data passing this inspection is passed on to the Online Computer System for further analysis and archival storage.

The coordination of the event data flow through the DAQ System is exercised by the Event Data Flow Control (EDFC) Subsystem. Its primary function is to keep event data flowing efficiently from front-end and trigger systems through the Event Builder and Online Processor Subsystems. The EDFC Subsystem continually monitors the condition of system elements affecting data flow, and influences the flow so as to eliminate pile-ups and bottlenecks. Triggers are inhibited due to impending buffer overflows in the readout crates and in the Event Builder Subsystem via "Inhibit Trigger" messages originating in the DAQ CPU modules and Event Builder Subsystem, respectively.

The DAQ Control/Monitoring Network monitors the integrity of the event data stream, provides alarms, error handling, calibration and monitoring services, controls the various processors and modules in the system and delivers events to consumer processes for miscellaneous on-line monitoring purposes.

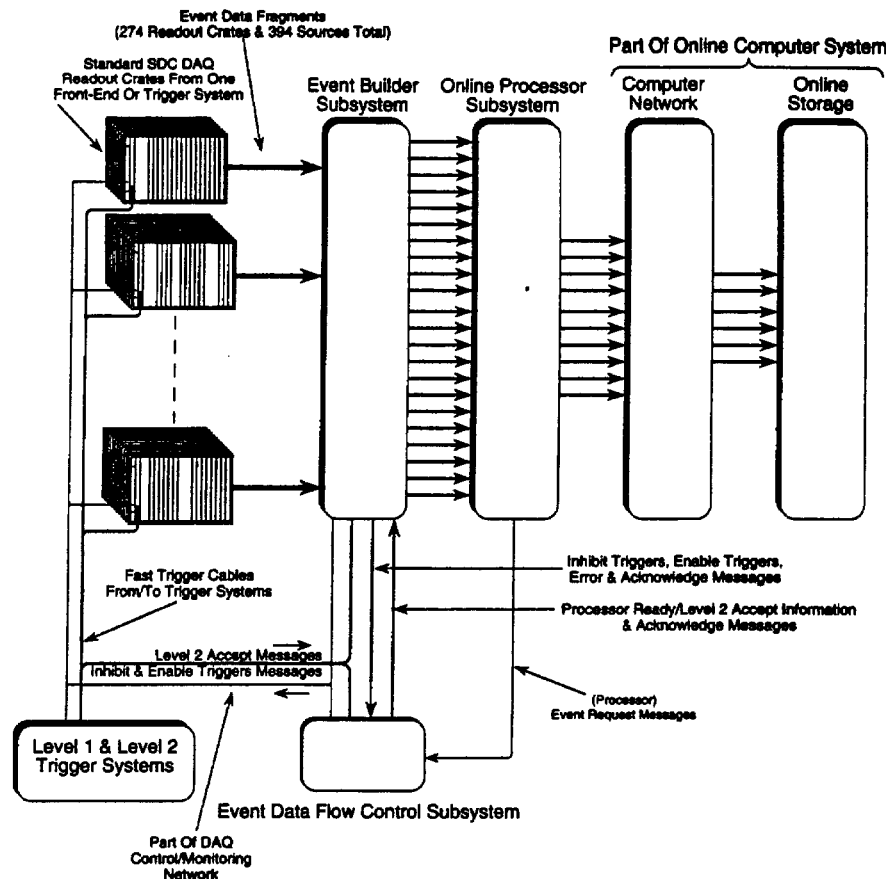


FIG. 8-46. Major components of the SDC data acquisition system.

The components of the DAQ System communicate with each other over data links of varying speeds. High-speed links carry data and messages among the front-end and trigger system crates, the Event Builder Subsystem and the Online Processor Subsystem. The Event Data Flow Control Subsystem interacts with the Trigger System, Event Builder and Online Processor Subsystems again over separate dedicated high-speed links. The DAQ Control/Monitoring Network is a low-speed Local Area Network that encompasses the front-end and trigger crates, DAQ CPUs, various workstations and embedded controllers that are distributed throughout the DAQ System.

A description of event data flow and its control is as follows (see Fig. 8-46). A Level 2 Accept message is sent by the Level 2 Trigger processor to all readout crates and to the Event Data Flow Control Subsystem via high-speed fiber optic cables. Event requests are sent to the Event Data Flow Control Subsystem by the Online Processor Subsystem whenever a processor is ready to process an event. The Event Data Flow Control Subsystem transmits both processor destination and Level 2 Accept information to the Event Builder Subsystem prior to the receipt of the first event data fragments resulting from the current Level 2 Accept.

Upon receipt of a Level 2 Accept message, the DAQ CPU modules within the crates initiate the readout of current event data fragments in one of two modes. In most crates, the CPU module gathers all current event data fragments from all front-end or trigger system modules in its crate. If the crate contains front-end modules that transmit their event data fragments directly to the Event Builder Subsystem, the CPU modules informs these modules of the Level 2 Accept message, thereby enabling them to transmit the fragments.

The CPU modules that gather event data fragments transmit the fragments to the Event Builder Subsystem. The Event Builder Subsystem receives and performs header and data integrity checks on all incoming event data fragments and then builds an event onto one its output ports where it is transmitted

to the processor in the Online Processor Subsystem that had requested an event. If no processors are ready to accept events when a Level 2 Accept Trigger is received, the Event Builder Subsystem simply buffers (temporarily stores) the events. If the stored data nears the capacity of the buffers, the Event Builder Subsystem inhibits triggers via the Event Data Flow Control Subsystem. Shown in Fig. 8-46 but not part of the DAQ System are the components of the Online Computing System used to transmit and permanently store data for further analysis. A description of the operation of these components is given in the Online Computing System chapter, Chapter 9. Additional details of event data flow and its control are given in Section 8.9.6.

8.9.3. Partitioning and stand-alone operations

The DAQ System architecture provides a number of ways of operating front-end systems in stand-alone mode or of partitioning the DAQ System itself. These include stand-alone operation using workstations connected through the DAQ Control/Monitoring Network to front-end and trigger readout crates, dynamic segment selection, partitioning the Event Builder Subsystem by front-end system, as well as partitioning the system at the individual readout crate level.

The simplest way of partitioning is to put an entire front-end system under the control of its Local Trigger Control Crate (referred to as the Level 1 Distribution Crate in Section 8.2). During initial commissioning, the DAQ CPUs of a given front-end system collect and buffer events which are then read out slowly via the DAQ Control/Monitoring Network. During the second phase of commissioning, the Event Builder Subsystem is logically divided by front-end system such that it functions as several independent event builders, one per front-end system, thereby providing each front-end system with independent readout and operation at full data acquisition speeds. When operating in a stand-alone mode, the Local Trigger Control Crate fans out the trigger signals to the front-end system. Thus, simply by telling the Local Trigger Control Crate to go into stand-alone mode, ignoring global triggers but fanning out an external trigger supplied by it to the entire front-end system, one can have an entire front-end system running coherently and independently in stand-alone mode.

Dynamic segment selection allows subsets of the data collection segments within a given partition to be read out in response to different trigger types. The trigger type is used to select which segments are read, so the subset of data read out can change from event to event. However, the correspondence between trigger types and data segments cannot change during a run. Dynamic segment selection can be implemented either by sending Level 2 Reject messages to those crates which are not part of the readout for that event, or by discarding the unwanted segments in the Event Builder Subsystem or Online Processor Subsystem if sufficient bandwidth is available in the data collection network and Event Builder Subsystem.

The Event Builder Subsystem routes data packets by destination, so many partitions can use the same switching network without interference. To improve bandwidth, the switching network can also be configured (through software) to operate as a group of independent switches. And, if complete isolation is required, the switching network can be divided into smaller physical switches.

Finally, each separate front-end readout crate can be put in its own partition simply by telling the Trigger/Gating module in that crate to go into stand-alone mode and by generating a trigger signal at the "external trigger" input to the Trigger/Gating module. Normally this would be for standard calibration runs, with the stand-alone triggers already connected to the Trigger/Gating module. These stand-alone triggers would become effective as soon as the Trigger/Gating module was commanded by the DAQ System to go into stand-alone mode. In this mode of operation, the DAQ CPUs would do the data acquisition for the crate, performing local processing as required, and transmitting the data (or summaries of the data) along its network connections to a front-end system workstation for logging or further processing.

8.9.4. Components of the data acquisition system

The following sections describe the various components of the DAQ System along with their proposed implementation. Some of the details of implementation may change with advances in technology, but the overall philosophy of using industry standards and commercial components, wherever possible, will remain the same.

Front-End and trigger system readout crates, modules, designs and links

The interfaces between the DAQ System and the front-end and trigger systems are distributed within the 274 readout crates. In order to promote uniformity of these interfaces, the purview of the DAQ System includes the design and production of the basic readout crate (the SDC Standard Crate) and two important interface modules that reside in each crate, the DAQ CPU module and the Crate Adapter/Interface module. It also includes the design and production of the physical and logical attributes of the various data links that carry data to/from and within the crates. The DAQ System provides two additional standard designs to be used in the implementation of front-end and trigger systems. These are the crate bus slave interface and the event data readout port interface. The position of the latter is shown in Fig. 8-47; additional descriptions of both components occur later in this section.

Figure 8-47 illustrates the two configurations of SDC standard DAQ crates. The top crate houses various front-end or trigger system modules together with a DAQ CPU module and a Trigger Gating module. In addition to these modules, the bottom crate also houses one or more Data Collection (buffer) modules which receive event data from on-detector front-end components. The Switched Capacitor Array Calorimeter/Shower Maximum detector and Outer Tracking Straw Tube readout crates use a configuration similar to that of the bottom crate. In this configuration, event data are not read out via the crate's bus but are transmitted directly from the Data Collection modules. All other front-end readout crates and all trigger system readout crates are in the configuration of the top crate; event data are read out via the crate's bus by the DAQ CPU and transmitted to the Event Builder Subsystem via the Crate Adapter/Interface module. For those readout crates with particularly low data rates (e.g., Level 1 and Level 2 Trigger System crates), a crate vertical interconnect could be used in place of a separate DAQ CPU module in each crate.

The Trigger Gating module in each readout crate is the interface between the crate and the Level 1 and Level 2 Trigger Systems. Trigger Gating modules are supplied and supported by the Trigger System implementers. These modules receive Level 1 and Level 2 Accept/Reject triggers and other fast signals from the Trigger System and transmit Inhibit Triggers messages to the Trigger System. Inhibit Triggers messages indicate that local conditions in the crate (e.g., buffers becoming full) require that triggers and event data flow be temporarily inhibited.

Powered crate and cooling unit

The SDC-standard readout Crate and Cooling Unit includes the crate, backplane, power supplies (2000W, +5V, +12V and -12V), air/water heat exchanger, fan unit, I/O unit, and the power distribution block (with the high-current connectors and connections to the backplane). This unit will be designed and constructed using industry standard crates, backplanes and power supplies. The decision to use 400 Hz or conventional 115 V, 60 Hz power will be made at a later date. Eurocard standard packaging of 9U \times 400 millimeters with VMEbus protocol is the present choice, but it is expected that this choice will be re-examined in the next few years as the system design is finalized. Power supplies will not be an integral part of a crate, but will be mounted behind the crate so as to allow uninhibited access to rear-connected cables.

A power distribution block consisting of high-current connectors connected to the VMEbus backplane is provided to improve power distribution as well as to eliminate many of the power supply cabling problems associated with VMEbus crates. A VMEbus I/O (input/output) Unit (auxiliary I/O card cage) is also provided to allow users a simple, SDC-standard method for connecting I/O cables to the rear of a module.




FIG. 8-47. Event data readout components and DAQ control/monitoring network interface.




FIG. 8-48. Physical layout of the crate DAQ adapter/interface module.

The air/water heat exchanger and fan unit, in conjunction with the cooling assembly of the Electronic Rack Protection and Cooling Subsystem (further description later) provides required cooling for both the VMEbus cards located in the VMEbus crate and the associated power supplies.

DAQ CPU modules

Since many of the processes associated with the collection, formatting and transmission of data are common to all front-end crates, and since these processes are intimately associated with the DAQ System, the DAQ System will provide a DAQ CPU module which is to be located in each readout crate. This will be a commercially available module which will contain a Local Area Network port to the DAQ Control/Monitoring Network.

Crate adapter/interface modules

As shown in Fig. 8-48, Crate Adapter/Interface modules physically adapt the smaller DAQ CPU modules so that they can be mounted in the larger crates. They also contain electronics for transmitting event data fragments to the Event Builder Subsystem, and for "extending" the DAQ Control/Monitoring Network to on-detector components not housed in SDC Standard DAQ crates (the Remote Control/Monitoring Link Port in Figs. 8-48 and 8-47).

Bus slave interface design

The Data Acquisition System will provide a complete design package, including specifications, timing diagrams and a behavioral model of the readout crate's bus slave interface. It is mandated that all front-end and trigger system modules adhere to these specifications. Assistance will be provided by DAQ System personnel to ensure front-end and trigger system module compliance to these specifications.

The bus slave interface includes logic to receive interrupts from the DAQ CPU such that, for example, the software protocol for informing modules of a Level 2 Trigger Accept is identical for all front-end and trigger system modules.

Event data readout port interface design

The Data Acquisition System will provide a complete design package, including specifications, timing diagrams and a behavioral model of the event data readout port interface. It is mandated that all front-end modules that transmit event data directly to the Event Builder Subsystem adhere to these specifications. Assistance will be provided by DAQ System personnel to ensure front-end and trigger system module compliance to these specifications.

Event data and event data flow control links

All event data and event data flow control links will use high-speed 1 Gbit/s fiber optic links. A link includes the driver and receiver logic and optics, connectors and cable. Excluding event ID and protocol header and error code overheads and a bandwidth safety margin, the transmission of real event data at rates of up to 60 Mbyte/s can be achieved on a link. The industrial market is driving down the cost of these links to a point where it will be a cost benefit to the SDC to standardize on these links not only for transmission of event data and fast control in the DAQ System but also for similar purposes in the trigger and other systems.

Event builder subsystem

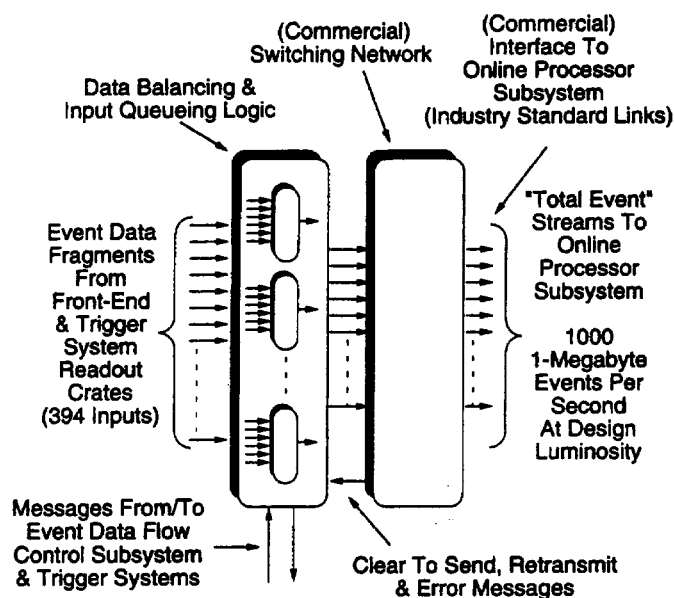
Event builders are used in large physics detectors to read out event data fragments from many sources and to collect these fragments into a data array representing a single physics event. Event builders have evolved from simple single-channel funnels through a minicomputer bus, to groups of parallel channels (each with their own funnel or bottleneck characteristics) feeding arrays (farms) of processors. As experiments have become larger and larger, the need for increased parallelism for higher throughput has grown. The severe conditions at the SDC (beam crossing time of 16 ns and very high trigger rates) impose the requirement for a totally parallel event builder with higher event data throughput rates than have ever been achieved to date.

The major specifications for the Event Builder Subsystem for the SDC detector include:

- Receive 394 streams (cf. Table 8-9) of event data fragments from front-end and trigger system sources with each stream having a bandwidth of 1 Gbit/s.
- Accommodate a throughput of typically 1 Gbyte/s scalable up to 10 Gbyte/s.
- Build events by combining the incoming event data fragments into a complete event record.
- Transmit complete event records into an online array of processors (the Online Processor Subsystem).
- In a modular way, accommodate expansion or contraction in the number of input and output channels and in the throughput of the subsystem.
- Connect to the Online Computer System via the DAQ Control/Monitoring Network.

The Event Builder Subsystem, as shown in Fig. 8-49, is based on a commercial switching network similar in nature to networks presently being used or being developed for the telecommunications industry (for high-speed parallel switching) and for the computer industry (for interconnecting computers and peripherals at very high-speed). This subsystem consists of Data Balancing and Input Queueing Logic

and a commercial Switching Network. The Data Balancing and Input Queueing Logic is used to receive incoming event data fragments from the 394 DAQ CPUs and event data buffer modules, to balance the load (data volume) over several events on any input of the Switching Network, and to queue equal-sized packets for transmission through the Switching Network. The queueing insures efficient high-throughput transmission of data through the switch. The Switching Network is used to build event fragments into a single event stream on any one output of the network. Details of this event-building operation are given in a later section. Both the Data Balancing and Input Queueing Logic and the Switching Network are scalable to accommodate future system growth.



(DAQ Control/Monitoring Network connections not shown)

FIG. 8-49. Event builder subsystem block diagram.

The Event Builder Subsystem receives Processor Ready/Level 2 Accept messages from the Event Data Flow Control Subsystem (see Fig. 8-46) indicating that a destination processor (of the Online Processor Subsystem/Level 3 Trigger) is ready to receive a particular event. The Event Builder Subsystem also uses the Level 2 Accept messages to perform integrity checks on the received event data fragments and informs the Event Data Flow Control Subsystem of any errors. The Event Builder Subsystem sends messages to the Event Data Flow Control Subsystem to inhibit triggers when its input buffers are becoming full, to re-enable triggers when its input buffers have sufficient capacity to accept events, to acknowledge receipt of incoming messages and to indicate errors. Input buffers become full if either triggers occur at a higher rate than can be accommodated by the Online Processor Subsystem or if certain errors occur (e.g., Switching Network malfunction).

The Event Builder Subsystem connects to front-end and trigger system readout crates and to the Event Data Flow Control Subsystem via industry-standard high-speed fiber optic serial links. Industry-standard high-speed links are also used to connect the Event Builder Subsystem to the Online Processor Subsystem.

Data balancing/input queueing logic

Switching networks are most efficient when the amount of data fed into each of their inputs is balanced over several events and when data packets of equal size are transmitted synchronously through the network. The Data Balancing/Input Queueing Logic of the Event Builder Subsystem performs these functions. In detail, this logic:

- Receives event data fragments, varying in size from a few hundred bytes to several kilobytes, from the 394 front-end and trigger system readout crate sources.
- Multiplexes data where appropriate (*e.g.*, small event data fragments) and forms fixed-sized packets of event data for transmission through the Switching Network such that each input of the Switching Network has roughly equal (balanced) amounts of data averaged over time.
- Time orders and synchronizes (queues) the transmission of the packets through the Switching Network such that no two packets have the same Switching Network output as their destination during any single packet transmission time.
- Communicates with the Event Data Flow Control Subsystem over industry-standard high-speed fiber optic links for the control of event data flow and error checking.
- Inhibits triggers when its input buffers are becoming filled (or there are fatal errors) by sending an "Inhibit Triggers" message to the Event Data Flow Control Subsystem (and onto the Trigger System).
- Performs data integrity checks on incoming data Level 2 Accept information from the Trigger System and processor ID information from the Online Processor Subsystem).
- Logs errors and keeps statistics on system dead time, errors *etc.*
- Communicates with the DAQ Control/Monitoring Network (not shown in Fig. 8-49).

The input interface of the Data Balancing/Input Queueing Logic includes fiber optic receivers and serial-to-parallel encoder logic. An input module will contain at least eight inputs; thus 128 to 152 inputs ($8 \times 16 - 19$ slots) can be housed in a crate. Three crates are therefore needed to receive the event data fragment on the 394 fiber-optic channels from front-end and trigger system readout crates.

The relationship between the Input Queueing Logic and the Switching Network is shown in Fig. 8-50.

Switching network

Both the telecommunications industry and the computing industry have been developing switching networks to increase system throughput and interconnectivity of their many nodes. Companies such as Ancor Communications [52], Network Systems [53] and UltraNet [54] have been providing switching networks, of varying size.




FIG. 8-50. Event data packet input queueing.

The Switching Network in the Event Builder Subsystem is a commercial switching network operating in a synchronous mode that can:

- Sustain throughputs of up to 1 Gbyte/s.
- Operate in non-blocking mode (an independent path from every input to every output).
- Expand in a modular way both in the number of inputs and in throughput,
- Optionally interface to commercially-available large processor arrays (Online Processor Subsystem) via several industry-standard links and protocols (*e.g.*, VMEbus, Fibre Channel [55], HiPPI [56], Serial HiPPI [57] and SONET [58,59]).
- Work with commercial operating system procedures *e.g.*, UNIX socket-type calls) or in a more efficient synchronous mode.

A prototype project which used input queueing and a switching network for parallel event building has been successfully demonstrated at Fermilab [60,61,62,63,64]. The project showed that multi-Gigabyte per second event building can be achieved in a straight-forward design by using switching networks. The project consisted of an 8-input, 8-output event builder (subsystem), an event data flow control (subsystem), eight data sources (test transmitter modules) and eight destinations (emulators of Level 3 Trigger processors). The event builder consisted of an input queueing network and a switch consisting of eight 64×64 crossbar ICs mounted on a special backplane. The project was simulated extensively and the above hardware/firmware was built and tested successfully. The event builder was packaged in a VMEbus crate and expanded to a 64-input, 64-output event builder with a 1.28 Gbyte/s aggregate throughput within a single crate. The prototype system was capable of expanding further to much larger numbers of inputs and outputs using additional crates. Events averaging 256 Kbytes of varying total event and event fragment sizes were inputted into the system's eight inputs at roughly 20 Mbyte/s per input including overheads, first with one, then two, four and eight outputs. Each doubling of the number of outputs doubled the event builder's throughput from 75 to 150, 300 and 600 events per second, respectively, thus proving out the parallel event building concept using a switch-based event builder.

Figure 8-51 shows an example of an operating mode of the Switching Network—i.e., the barrel-shift mode. For demonstration purposes, the figure assumes four event fragments per event each of equal size and each constituting a fixed-length packet. The three rectangles preceding the switch each represent a single fixed-length packet delay in the Input Queueing Logic. During the first rotation of the switch, each input is connected to its corresponding output. The first component of event 1, event data fragment 1A is transmitted to output 1 and null packets are transmitted to outputs 2, 3 and 4. During the next rotation of the switch, inputs 1, 2, 3 and 4 are connected to outputs 4, 1, 2 and 3, respectively, such that the second component of event 1, event fragment 1B, is transmitted to output 1 and the first component of event 2, event fragment 2A, is transmitted to output 4. As the switch rotates through its successive configurations, streams of "built" events are outputted in parallel.

Various operating modes of the Switching Network, both synchronous and asynchronous, are possible. Multiple event data fragments from several events can be combined into single fixed-length packets and concatenated into single events in the Online Processor Subsystem for increased throughput. The Event Builder Subsystem can operate either in an asynchronous mode, sending event fragments through the switching network as they arrive and as a destination is ready (UNIX socket-like mode), or in a totally synchronous mode, as in the above example, for increased throughput. In this latter mode, the Clear To Send message (or signal in this case) shown in Fig. 8-49 indicates to the Input Queueing Logic that the previous synchronous packets have been transmitted through the switch and the switch is reconfigured and ready to receive and transmit the next set of packets.

When a switching network is used for general processor interconnection, there are delays in configuring the switch. Blocking occurs when two processors transmit packets to the same destination, in which case the switch controller must arbitrate. When a switching network is used as an event builder, the switch control can be simplified or even eliminated. If the switch follows a fixed rotation pattern, the queueing

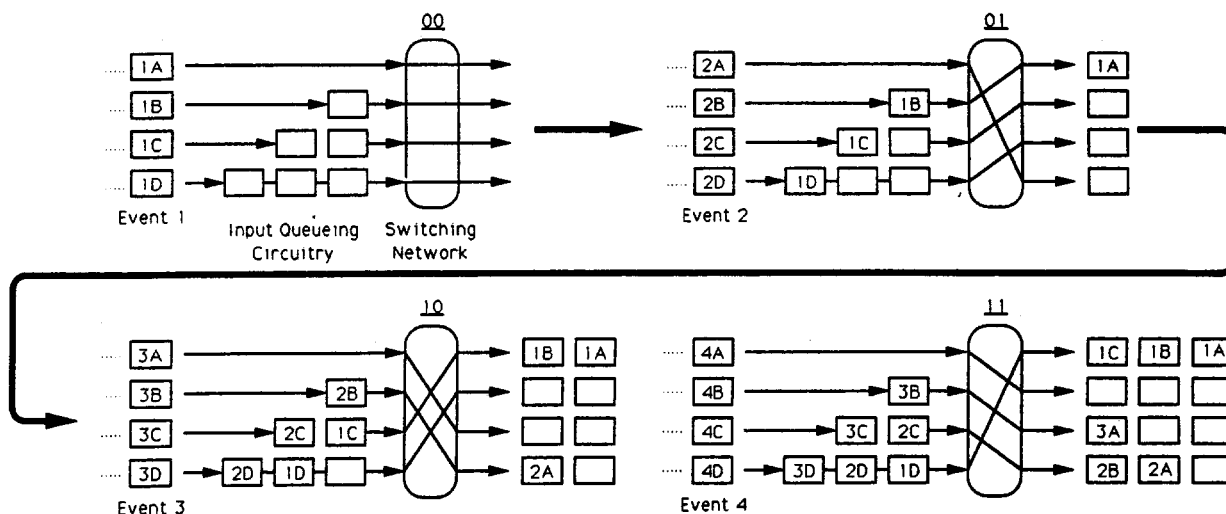


FIG. 8-51. Barrel shift switching network operation.

system will always know the next output connection and can sequence packets in the proper order, all with different destinations to avoid blocking or transmission delays.

The number of Switching Network outputs needed to sustain a given output bandwidth depends on the bit rate of each output and the overheads associated with the transmission of data. Assuming an industry-standard protocol and physical link, such as Fibre Channel, normal overheads include those associated with the protocol, data bit encoding, physics headers, load imbalances on average between Switching Network outputs (different event sizes), and throughput safety margin. All these overheads combined could result in actual event data transfer rates 1/3 to 1/2 of the theoretical bandwidth of the links. Thus, to achieve an initial actual event data bandwidth through the switching network of 1 Gbyte/s at a Level 2 trigger rate of 10^3 per second, the initial implementation of the Switching Network will use a 32-input, 32-output switching network with 1 Gbit/s links. This will be scalable to larger numbers of inputs and/or outputs as necessary in future upgrades.

The method of operation for the Event Builder Subsystem described in this section is synchronous in that packets are transmitted synchronously through the switch, each going to a different output, with the Event Data Flow Control Subsystem functioning as a centralized control unit. In addition, the Data Balancing/Input Queueing logic interconnects readout crate sources to the Switching Network. An asynchronous method of operation of the Event Builder Subsystem, representing more of a general network, is currently under study. In this method, a larger Switching Network is logically divided into three logical switches such that readout crate sources connect directly to the first logical switch. Single input, single output intelligent data concentrators then connect that logical switch such that the Switching Network itself is used to multiplex (concentrate) readout crate sources. The outputs of this set of intelligent data concentrators connect to a second logical switch. The outputs of the second logical switch connect to another set of single input, single output intelligent data concentrators whose outputs, in turn, connect through a final logical switch to the Online Processor Subsystem. The first and second set of single input, single output intelligent data concentrators are used to buffer and concatenate partial and full events, respectively. In this asynchronous mode of operation, there is no centralized control (no Event Data Flow Control Subsystem). All links are bidirectional (two cables). Event requests from processors in the Online Processor Subsystem are transmitted through the Switching Network to data concentrators (data balancing/input queueing devices) which, in turn, transmit these requests to the readout crates.

Event data flow control subsystem

The coordination of the event data flow through the DAQ System is exercised by the Event Data Flow Control (EDFC) Subsystem. Its primary function is to keep event data flowing efficiently from front-end and trigger system readout crates through the Event Builder Subsystem and to the Online Processor Subsystem. Communications between the Online Processor Subsystem and the Online Computer System are used to transfer selected Level 3 events to online archival storage. The EDFC Subsystem has been designed to exercise its event data flow control using a relatively simple but secure protocol while processing events at rates of up to 10,000 per second.

As shown in Fig. 8-46, the EDFC Subsystem exchanges messages with the Level 2 Trigger System, the Event Builder Subsystem and the Online Processor Subsystem. These messages contain information regarding the availability of data destined for the DAQ system, and the present capability of the Event Builder and the Online Processor system to absorb and process these data.

The trigger system, in effect, "creates" event data sets by issuing Level 2 Accept triggers to the readout crates of front-end and trigger systems. The DAQ CPU module in each crate then either enables front-end modules to send their event data fragments or reads the event data from its front-end and trigger system modules, and formats and sends an event data fragment to the Event Builder Subsystem's Input Queueing/Data Balancing Logic. Coincident with each accepted Level 2 trigger, the trigger system sends a Level 2 Accept message to the EDFC Subsystem. Included in this message is the identification (serial number) of the event. The message also includes information regarding the type of trigger processing used for that event, which may imply the class of online processor to which it will be sent.

The EDFC Subsystem also receives processor Event Request Messages containing available processor ID (destination) information and, optionally, trigger type information from the Online Processor Subsystem via standard serial Local Area Network links. These messages enable the EDFC Subsystem to maintain awareness of the availability of processors in the Online Processor farm and thereby to assign idle processors to incoming events.

In response to a Level 2 Accept message, assuming the EDFC Subsystem is aware of an available processor, the EDFC Subsystem sends a Processor Ready/Level 2 Accept message to the Event Builder Subsystem. This message contains the event identification, which enables the Event Builder Subsystem to identify and properly assemble the incoming event data fragments into a complete event. It also contains an identification of the processor that the EDFC Subsystem has assigned to process that particular event. In response to this message, the Event Builder Subsystem will forward the set of incoming data for that event to that processor by routing it through the appropriate Switching Network output link. If no processors are ready to accept an event when a Level 2 Accept message arrives, the EDFC Subsystem queues the Level 2 Accept message and forwards the Level 2 Accept information to the Event Builder Subsystem. The Event Builder Subsystem's Input Queueing logic buffers the resulting incoming event until receipt of a corresponding Processor Ready message whereupon the buffered message is transmitted through the switch to the Online Processor Subsystem.

The EDFC also receives various types of messages from the Event Builder Subsystem. The message types include Inhibit Triggers, Enable Triggers, error condition reports and acknowledgements. An Inhibit Trigger message, for example, may be sent because some data buffer within the Event Builder Subsystem is becoming full or, for some other reason, it cannot accept more event data at this time. To disable the flow of event data, the EDFC Subsystem will send an Inhibit Triggers message to the trigger system. When the "becoming full" condition has cleared, an Enable Trigger message will be sent from the Event Builder Subsystem to the EDFC Subsystem and from EDFC Subsystem to the trigger system.

The EDFC Subsystem also connects to the Online Computing System via the DAQ Control/Monitoring Network. Various logging and statistical functions can be performed in the EDFC Subsystem if desired.

All message links to and from the EDFC Subsystem are serial high-speed fiber optic links.

Online processor subsystem (Level 3 trigger)

The Online Processor Subsystem consists of a farm of approximately 1000 individual physically small, but powerful processors, each having a processing power equivalent to that of 100 VAX 11/780 processors. In normal operation, each processor receives complete ensembles of event data from the Event Builder Subsystem and processes the data with appropriate software. In part, the processing will determine whether the event possessed certain physics qualities that make it worth saving in an archival storage. This is the Level 3 triggering operation. Events that pass the Level 3 criteria are output to the online data storage facility.

The Online Processor Subsystem consists of a large number (farm) of computers (~ 1000 initially and 5000 at full luminosity) operating in parallel. The architecture of the processor farm is a simple MIMD (multiple instruction/multiple data) computer architecture. In a MIMD architecture, each processor operates independently of all the other processors on an independent set of data—a separate event—with a high-speed input and output port for each processor unit. High-speed intercommunication between processor units is not necessary.

The result of the data processing is a decision on whether the event should be accepted and recorded on the permanent storage medium. This third stage of the trigger differs from the two earlier stages in that it is purely a software decision, with the trigger algorithms written in high-level programming languages. Both the raw data and the processed data of accepted events are written into a large data storage system. Most importantly, the software environment in this Level 3 Trigger processor farm must support the full offline computing environment, including operating system calls and data base accesses, so as to allow the appropriate portions of offline reconstruction code to be used in the final stages of event selection.

Processor farm

The requirement for total processing power of the Online Processor Subsystem is estimated to be 10^5 MIPS (million instructions per second, in VAX 11/780 units) at the expected initial input rate of 10^3 events per second. This requirement has been derived from current experimental experiences at the CDF detector at Fermilab and extrapolation to the anticipated data analysis needs of the SDC detector. The size of an SDC event is estimated at 1 Mbyte with a typical multiplicity of 300 in comparison with 100 Kbyte and 50, respectively, from the CDF detector. The required CPU power (MIPS-seconds) for the analysis is assumed to be between 100 and 1000 MIPS-seconds compared to 20 MIPS-seconds for CDF. Given the initial expected typical input rate to the Online Processing Subsystem of 1000 events per second, this results in a total processing power requirement of 10^5 MIPS at design luminosity of $10^{33} \text{cm}^{-2} \text{s}^{-1}$, with increased requirements at higher luminosities. The event data will be fed to the processor farm through multiple input channels in parallel from the Event Builder Subsystem.

The conceptual block diagram of the processor farm is shown in Fig. 8-52. The detailed architecture of the processor farm will be designed to satisfy the following key design parameters: CPU power, required CPU time for the realistic event data, number of input channels, and possible bandwidth of the channel. In Fig. 8-52, parameters of the conceptual design are indicated. The computing power assumed is 100 MIPS per processor, which implies a processing time of 1 to 10 seconds for a typical event (based on a typical processing load of 100 to 1000 MIP-seconds per event). The figure shows 32 input channels each with 32 processors per channel. At an incoming event rate of 1000 events per second with 1 Mbyte per event (or 1 Gbyte/s aggregate) and assuming a 50% factor in actual data input rates due to overheads, etc., each input channel must have a capacity of roughly 62 ($[1000/32]/0.50$) Mbyte/s. The initial design of the processor farm will use 1 Gbit/s input data links. Reasons for this choice are given below.

Input and output data links

The input data links to the processor farm from the Event Builder Subsystem must be designed with sufficient bandwidth to accommodate the event data generated by the expected trigger rates. The expected overall event data rate, assuming a Level 2 trigger rate of 1 kHz and 1 Gbyte of data per event, is 1 Gbyte/s. However, the system must be designed to accommodate a 10 kHz Level 2 trigger rate, which results in an overall event data rate of 10 Gbyte/s. This total data rate must be apportioned over an appropriate number of individual data links to mesh with the time to process an individual event and the number of processors in a branch. For example, a match to the design data input rate capability of the processor farm of 10 Gbyte/s (*cf.* Fig. 8-45) can be achieved using approximately 160 1 Gbit/s input links. This number is calculated using a value for overhead of 50%, as earlier explained in the discussion of the Switching Network. Expected technology improvements will undoubtedly provide industry standard links operating at 1 Gbit/s by the mid-1990s.

The output of the processor farm must transfer event data blocks that have been processed by the processors, and transfer the blocks to the online archival storage facility. Because of the filtering of events by the Level 3 filtering, the required data rate of the output data links is assumed to be reduced by a factor of 10 to 100 from the links at the input to the processor farm.

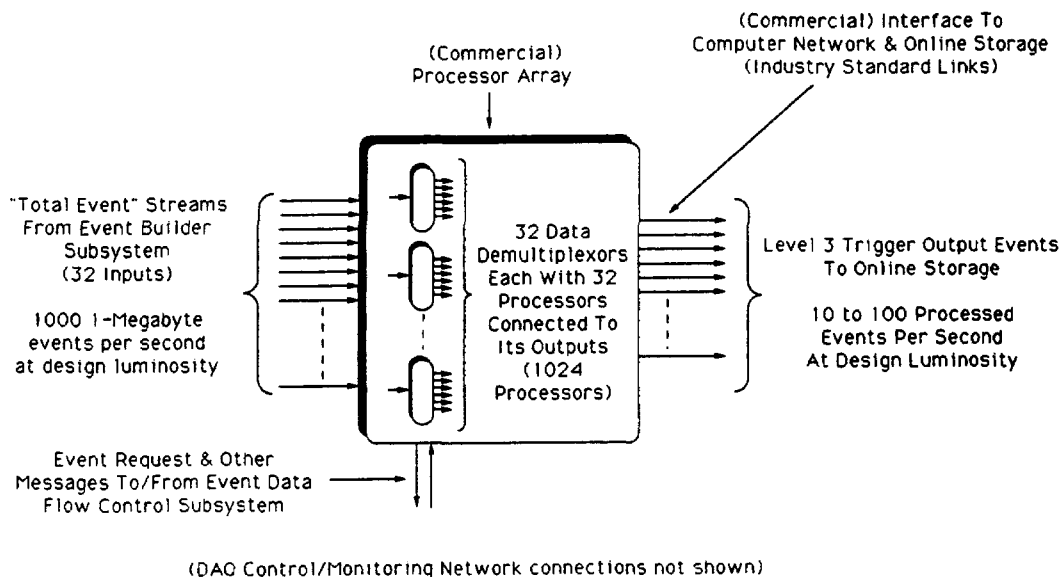


FIG. 8-52. Online processor subsystem processor farm conceptual block diagram.

The required overall output bandwidth will therefore be roughly 100 Mbyte/s (100 events per second) at design luminosity. It would be convenient if the implementation (protocol, physical layer, *etc.*) of the output data links was the same as that of the input data links, but this is probably not absolutely necessary if there are good reasons (such as compatibility with the online computing system) to use a separate protocol. If the implementation for both is the same, and both use 1 Gbit/s links, one tenth as many output links as input links will be required.

Within the communications community, several data link protocol and physical layer standards are either being implemented or being specified. Fibre Channel [55], a bit-serial version of High Performance Parallel Interface (HiPPI) [56], called Serial HiPPI [57] and SONET [58,59] are currently being implemented and a bit-serial version of Scalable Coherent Interface (SCI) [65] is currently being specified. All operate at 1 Gbit/s over fiber optic cable. It appears that Fibre Channel will be used extensively for processor and peripheral interconnections by the mid-1990s.

The final design of the processor farm will be done in the mid-1990s. These industrial data link standards are recommended to be used because of the reusability of the software system.

Several companies are developing large arrays of processors. One such company, IBM [66], is working with Ancor [52] to integrate Ancor's switching networks with IBM's large arrays of RISC processors. It is evident that both switching network and computer companies will continue to provide several industry-standard links and required communications software for interconnecting their respective products. VMEbus, Fibre Channel, HiPPI, Serial HiPPI and SONET will be included in these options. Therefore, the Event Builder Subsystem's Switching Network outputs will be able to be directly connected to commercially-available large processor arrays (Online Processor Subsystem) using one of these link standards.

Operating system and other software

The operating system of the Online Processor System will be fully compatible with that of the Offline Computing System to facilitate the software exchange between them. The choice must be subject to change following industrial trends. The Data Acquisition group will develop a software package to manage the various processors in the farm, including downloading, delivering events to the processors, monitoring performance, debugging and detecting failures.

This software package will be based on software used in existing offline and online processor farms that have been in use for more than five years at Fermilab and elsewhere. Ordinarily each processor will run an identical program on different events, although the architecture also allows using different programs in different processors (*e.g.*, for different trigger types). The software systems track the performance of each processor in the system, delivering real-time information on processor utilization and error conditions. The most likely errors to occur are user software bugs that cause the trigger algorithms to crash on rare events. These errors will cause the affected processor to be downloaded with a fresh copy of the program and restarted. Repeated hardware failures in a single processor will cause that processor to be removed from service.

The system will be designed by the mid-1990s for prototyping and to develop the software for the system. However, the actual acquisition of the processors will be made as late as possible since the cost of these processors is continuing to fall rapidly. In fact, only half of the full on-line processor subsystem (5×10^4 MIPS) is costed in the SDC baseline, since we assume that the accelerator will not be performing at full design luminosity at initial turn on. The acquisition of the remainder of the processors will be made as a detector upgrade as the machine achieves full design luminosity, with the cost of the processors minimized the longer their acquisition can be delayed.

Online storage subsystem

The Online Storage Subsystem, although not part of the DAQ System, is described here for completeness. This system has to accept the data at the rate of 100 Mbyte/s at the typical event rate. The required data recording rate is 100 Mbyte/s. A conceptual design with currently available technology is as follows. The data recording rate of a 19 millimeter magnetic tape is 30 Mbyte/s. The capacity of the tape is 100 Gbyte per cartridge. The cartridge tape is a standard tape for a TV broadcasting station. A library system for this cartridge with a capacity of 27 Terabytes will soon be commercially available. Based on this system, the output network of the Online Processor Subsystem is organized to converge into three branches each with the controller of the tape system. Although the bandwidth of 30 Mbyte/s is current technology for the data recording, the final system can be designed with more advanced technology, as available. More details on the Online Storage Subsystem are discussed in the Online Computing System, Chapter 9.

Data acquisition system control/monitoring network

All Data Acquisition System non-event data control and monitoring functions are performed over the DAQ Control/Monitoring (Local Area) Network. This network connects directly to the Online Computer System and indirectly to the Ancillary Controls System via the Ancillary (Slow) Controls (Local Area) Network as shown in the example implementation in Fig. 8-53. Databases in various Online Computer System computers and the network routers segment the DAQ Control/Monitoring Network by front-end and trigger system. This segmentation allows parallelism during Data Acquisition System downloading, *etc.* for increased throughput and allows commissioning and debugging of each front-end or trigger system independent of any other system. The workstations provide local control locations and a path for slow event data readout during initial commissioning and also simplify system maintenance. An Ethernet/Thinwire Ethernet DAQ Control/Monitoring Network was costed but the "network of choice" at the time of production ordering would actually be used in the Data Acquisition System. Studies are needed to determine the correct network implementation for the proper amount of parallelism and bandwidth to meet, for example, system downloading time requirements. For illustration purposes only, Fig. 8-53 shows an Ethernet and Thinwire Ethernet implementation. The Local Area Network chosen at production ordering time will be that which meets system requirements and is readily available from industry at the time. Routers make possible initial commissioning of individual front-end and trigger systems without interference from other systems. These routers also provide a means for the workstations to access data from more than their particular system during partial integration tests. The workstations shown, one per front-end and trigger system, simplify commissioning and other functions. These workstations may become unnecessary as the system matures.

Both the Online Computer System and the Control Computer System have access to all data acquisition components through the DAQ Control/Monitoring Network. Details of the control/monitoring application programs provided by Data Acquisition System are given in Section 8.9.5.

Interfaces to the online computer and ancillary controls systems

Figure 8-53 also shows the network connections of DAQ Control/Monitoring Network to both the Online Computing and Ancillary Controls Systems. The same networks and software protocols are used for both Data Acquisition and Ancillary Controls control/monitoring functions. The software applications that provide the DAQ Control/Monitoring Network functionality can thus be run on any of the computers in the system. Moreover, the low-level software (for downloading or sending a command to a particular address of a particular crate) provided by the Ancillary Controls group will be used in the Data Acquisition Control/Monitoring applications. Similarly, the same user interface should be used for data acquisition, control, and online applications.

Connections to data acquisition system components

The DAQ Control/Monitoring Network connects to all DAQ System subsystems and readout crates except for the Electronics Rack Protection Subsystem, which is connected directly to the Ancillary (Slow) Controls LAN. The DAQ Control/Monitoring Network connects to front-end and trigger readout crates via the DAQ CPUs within these crates.

Connections to remote front-end and trigger components

Many pieces of electronics will not be mounted in crates, but rather on or very near the detector they are associated with. These pieces need communication facilities to provide for various control and monitoring functions. Remote Control/Monitoring links and repeaters are used for connecting these devices indirectly to the DAQ Control/Monitoring Network via circuitry on the Crate Adapter/Interface modules in nearby crates. To exercise control and monitoring, DAQ Control/Monitoring Network messages are received by DAQ CPU modules within the readout crates, re-formatted and then forwarded to the Crate Adapter/Interface modules. From here they are transmitted to the pieces of electronics through repeaters and over simple industry-standard serial links, such as IEEE P1197 (Test Scan). Figures 8-47 and 8-48 show




FIG. 8-53. DAQ control/monitoring network example implementation.

the port on the Crate Adapter/Interface module in SDC Standard DAQ readout crates used to interface to these links and repeaters.

Electronics rack protection and cooling subsystem

The Electronics Rack Protection Subsystem is used for both hardwired and computer-controlled protection of electronics in trigger, front-end and data acquisition racks. The subsystem will most likely be a closed air/water heat exchanger system. It is planned to provide for remote monitoring of parameters such as air speed, water dripping, smoke, water flow, humidity and power supply voltages. A hard-wired automatic power off will occur if, for example, smoke or insufficient air flow is detected.

Items protected with this subsystem include all Front-End and Trigger System readout crates, the Event Builder Subsystem and the Event Data Flow Control Subsystem. If the commercial Online Processor Subsystem is not well protected as delivered, it will also be protected (and monitored) by the Electronics Rack Protection Subsystem.

Initial plans are to avoid duplication of design effort by using a electronics rack protection subsystem developed by the SSC Laboratory for its systems.

For safety and system integrity reasons, the Electronics Rack Protection chassis within each rack is powered separately from other electronics in the rack. For the same reasons, control and monitoring of this subsystem is via the Auxiliary Controls System's Supervisory LAN rather than the DAQ Control/Monitoring Network.

8.9.5. System software

The software for the Data Acquisition System represents a complex interplay of systems and applications programs written by the Data Acquisition group, the Online Computing group, the Ancillary Controls group, and by front-end and trigger system physicists. In general, the Online and Ancillary Controls groups will provide system software, tools and templates to support applications programs. The Data Acquisition group will provide further tools and templates (particularly relating to real-time applications) as well as certain generic applications programs (such as start-up and run control for the Data Acquisition System and alarms and error logging applications). The front-end and trigger system physicists will write most of the actual front-end and trigger system-specific application programs for calibration, monitoring, and control, using the tools provided by the other three groups. The various software components that the Data Acquisition group will provide are described below.

Runtime user interface

The Runtime User Interface is the means through which users will interact with the system during physics data taking. Most of the applications software items included in the Runtime User Interface are listed below:

- Run Time Control
- Control of the Online Processor/Online Storage Subsystems
- Histogramming
- Event Displays
- Calibration Database
- Consumer Processes

Many of these tools (histograms, event displays, database interface, user graphic interface, hooks for user consumer processes) will be provided by the Online Computing group. The Data Acquisition group will provide additional tools specific to the real-time environment running closer to the front-end systems. In addition, the Data Acquisition group will be responsible for application programs that initialize, start and control the data acquisition, front-end and trigger systems (including booting, downloading, system checkout, performance monitoring, and error detection) as well as control of the Online Processor Subsystem.

Control, monitoring and downloading interface

Control, monitoring and downloading of software into the various processors distributed throughout the SDC electronics system are three essential functions of a Data Acquisition System. A partial list of subsystems needing control and/or monitoring applications software is given below:

- Power and Grounding
- Gas
- High-Voltage
- Environmental Monitoring
- Magnet Monitoring
- Calibration Pulsers
- Cryogenic Control System
- Accelerator Controls Interface
- Luminosity Monitors
- Manipulators

The low-level routines that actually interact with hardware components to perform these control functions will be provided by the Ancillary Controls group, while front-end and trigger system groups will write the high-level routines that users will interact with. The Data Acquisition group will provide both systems and application level support as described in the following subsections.

Software for supporting memories and intelligent processors

The Data Acquisition System will contain a wide variety of “embedded” processors that perform various functions requiring varying degrees of intelligence. The Data Acquisition System must provide services to these processors that enables them to be booted and downloaded, that allows them access to network and other I/O services, and that allows programs running in these processors to be interactively debugged. A user must be able to log in to these processors to run debugging or diagnostic code. Control capabilities must be present to start up the various processors, download them with appropriate code and data for the task at hand, detect errors, restart when necessary, etc.

Some functions provided by this applications software are:

- Downloading
- I/O (file) system services
- Terminal connections for debugging
- Network connections

Software for supporting non-event data acquisition

Functions supported by this class of software include:

- Detector parameter monitoring and logging
- Calibration data acquisition
- Data base access

Aside from the high-speed event data acquisition, the system must also support the collection and recording of other types of data not specifically associated with a single interaction. These include periodic monitoring of environmental, accelerator and detector parameters (temperature, high voltages, magnet currents, luminosity, etc.).

A variety of different frequencies for this slow data acquisition must be supported. Variances with normal conditions should be detected and reported. In addition, the system must support various types

of data acquisition, for calibration purposes, at a subsystem level where, for example, events are recorded without sparsification or pedestal subtraction, and must provide the processing power necessary to reduce this calibration data to physically meaningful quantities. Finally, the system must provide access to a data base that records the time history of both calibrations and other detector parameters.

The actual slow data acquisition will be done by low-level routines provided by the Ancillary Controls group, but the Data Acquisition group will need to provide standard slow-data acquisition processes and templates that will allow front-end and trigger system physicists to assemble such applications quickly and easily.

Software for recording unusual occurrences

Some of the reasons for requiring this class of software are as follows:

Given the complexity of the Data Acquisition System, faults and other error conditions must be detected automatically and quickly fixed. This will require a diagnostic connection to various electronic components, a set of test procedures that exercise system functions at start-up and periodically thereafter, monitoring procedures that check for errors in normal data taking conditions, and a comprehensive alarm and limit system which records alarms of varying priorities and brings those alarms requiring human intervention to the immediate attention of operators. This software provides these functions:

- Supervision remote control/monitoring links
- Execution of online diagnostic routines
- Supervision of alarms and limits throughout the system

Software for control and monitoring of adjustable devices

There will be a wide variety of devices on the SDC detector that need to be set, adjusted or manipulated remotely. A uniform software interface to these systems must be provided that allows the experimenters to set parameters, restore parameters to standard conditions, and record status conditions.

Most of this software is provided by the Ancillary Controls group, but the Data Acquisition group will provide software to present a standard user interface for these control and monitoring functions.

8.9.6. Event data flow, its control and trigger inhibiting

The following is an example of how the system operates (refer to Fig. 8-46). During system start up, the system is first reset (buffers empty and triggers and event requests disabled). Either Level 2 and Level 1 triggers or Online Processor Subsystem processor event requests can be enabled first for graceful system startup. When a Level 2 Accept message from the Level 2 Trigger System is received, the Event Data Flow Control Subsystem quickly matches this message to a Online Processor Subsystem processor event request and forwards a Processor Ready/Level 2 Accept Information message to the Event Builder Subsystem. This message not only contains event trigger type and other trigger information but a destination address (Switching Network output link and other processor ID information). Receipt of this message enables the Event Builder Subsystem's transmission of the particular event through its Switching Network and to the Online Processor Subsystem. When the Trigger/Gating module in a front-end and trigger system readout crate receives a Level 2 Accept message, it notifies the DAQ CPU. The DAQ CPU then reads its event data for the accepted event and transmits it to the Event Builder Subsystem via the Crate Adapter/Interface module's event data readout port. Alternatively, as in the case with Central Calorimeter Switch Capacitor Array and Shower Maximum detector data, the DAQ CPU informs Data Collection buffer modules within its crate to transmit the accepted event, via the module's event data readout ports, to the Event Builder Subsystem.

If Level 2 Accept messages start and continue to be received with no received Event Request messages (or no received Event Request messages for that type of trigger), the Event Data Flow Control Subsystem, after forwarding the Level 2 accept message to the Event Builder Subsystem, queues these messages until its accepted event trigger queue buffer starts becoming full. An Inhibit Triggers message is then sent to the

Trigger System to stop any further Level 2 Trigger System triggers. Total systems engineering ensures that triggers are inhibited before the accepted-event trigger queue buffer becomes full. When the first Event Request message that can be matched with a queued accepted event trigger is received, the Event Data Flow Control Subsystem sends an Enable Triggers message to the Trigger Systems—re-enabling triggers—and forwards a Processor Ready message to the Event Builder Subsystem. Event data fragments received at the inputs of the Event Builder Subsystem are buffered but not transmitted through the Switching Network until corresponding Processor Ready messages are received.

If Processor Ready/Level 2 Accept messages are received quicker than events can be transmitted through the Event Builder Subsystem, and the Event Builder Subsystem input buffers are becoming full, an Inhibit Triggers message is sent to the Event Data Flow Control Subsystem and forwarded by that subsystem to the trigger system to inhibit triggers. Again, total systems engineering ensures that triggers are inhibited before these buffers become full. The Event Builder Subsystem sends an Enable Triggers message when it has sufficient buffer space to accept additional event data fragments from readout crates.

A timeout in the Online Processor Subsystem is used to indicate too much time elapsed before a requested event was received. The cause of the timeout could be a system error or that no events or too few events of a certain trigger type have occurred.

Event Request messages received by the Event Data Flow Control Subsystem are queued and matched with corresponding events either on a trigger-type basis, on a first-come-first-serve basis, or on any other desired criteria. If event request by trigger type is supported, there will be an event request queue in the Event Data Flow Control Subsystem for each trigger type.

Because the Event Builder Subsystem receives both Level 2 Accept messages directly from the trigger systems and Processor Ready/Level 2 Accept Information messages from the Event Data Flow Control Subsystem, a high degree of data integrity checking (event ID, trigger type, etc.) in addition to transmission error checking can be performed in the Event Builder Subsystem. This data integrity checking can also be performed in the Online Processor Subsystem and offline as well.

8.9.7. Error detection and recovery

Special care has been taken to ensure that the control of the flow of event data through the system is very simple and errors encountered in this flow are both easy to detect and correct. The number of messages to sustain event data flow is held to a bare minimum to increase system throughput. Circular event data header, event data and event data flow control diagnostic buffers are included at various points in the system allow for simple error detection and recovery.

If single failures of critical parts do not disable the system then, because the architecture permits simple error detection and fast error correction (e.g., via diagnostic buffers throughout the system, access to error and statistical information in a simple manner over links totally independent of event data flow, etc.), the downtime of the system should be minimal.

The following paragraphs discuss features of the system's architecture and modes of operation which allow simple error detection and recovery.

Circular diagnostic buffers are implemented in software in each readout crate's DAQ CPU, in hardware in each Crate Adapter/Interface module, in each channel of the Event Builder Subsystem's Data Balancing/Input Queueing logic and in the input circuitry of the Online Processor Subsystem. These buffers record the header (source, destination, event IDs) and disposition (processed or discarded due to error) (and, under specific request, the data) of the last several incoming event fragment packets. A diagnostic task will occasionally read these buffers to determine if any long-term faults exist. The location of the faults can be determined by comparing the progress of packets through the system. Similar buffers are implemented in the Event Data Flow Control Subsystem.

By receiving Level 2 Accept messages, the Event Builder Subsystem can perform integrity checks on the received event data fragments and inform the error logging functions within the Event Data Flow Control Subsystem of these errors by transmitting an error message to that subsystem.

The simplest approach to error recovery is to flush and reset the data pipeline. This mode of operation does not imply any particular data format or system implementation. A checksum on each packet of data is sufficient to detect errors. Buffer overflows (although they should not occur), missing or extraneous data and routing errors all result in incomplete events or invalid checksums at the destination processor. If the error were caused by a corrupted event data fragment packet header (especially word count or destination information), all subsequent events may be misaligned. Flushing the pipeline removes these misaligned events, which may number several thousand.

The second and preferred approach is to treat each data packet independently and discard the bad packets. This is the approach taken at the lower level in most data networks. The system is reset only if the errors are continuous. Each packet must be "self-contained" with all necessary header information (event ID, packet ID, source, destination and checksum or CRC). Framing must be performed in hardware to allow re-synchronization on packet boundaries. If an error is detected, the entire packet is discarded. There is no attempt to forward partial packets. An error in a data packet for one event will not normally affect any subsequent packets or events.

Fixed packet lengths are not essential but may greatly improve the efficiency of the Switching Network. Fixed packet lengths also simplify buffer usage and overflow detection. If there is room for one complete packet in the buffer, the packet is accepted, otherwise it is discarded. Buffer management is more difficult for variable length packets where the word count information must be reliable. The disadvantage of fixed packet length is the overhead of the last packet which is half filled (on average). This effect may be minimized if the packet size is small.

Packets are created at the DAQ System readout crates. A readout crate may send more than one packet for each event, but the packets must be numbered (packet ID) and the last packet must also be labeled with an end flag. To distinguish missing packets from empty packets, each source must send at least one packet for each event, even if the packet contains no data. Except for framing errors, the packet is not checked until it arrives at the destination processor.

The destination processor or input circuitry of the Online Processor Subsystem checks for sequentially numbered packets from each source with the correct event ID and destination headers. If any packet in the event is missing or has an incorrect CRC, an error summary is sent to a diagnostic processor.

8.9.8. Stand-alone data acquisition systems

There is a need for stand-alone (or portable) data acquisition systems long before the complete Data Acquisition System is assembled. In fact, stand-alone data acquisition systems will be required for use in test beams and in lab tests of prototype and pre-production versions of front-end electronics and detector systems, as well as for testing of mass-produced equipment before installation in the detector. It is important that the stand-alone data acquisition systems use the same protocols and interfacing conventions as the full-scale Data Acquisition System in order that the modules under test experience the same environment as they will encounter in their ultimate locations in the SDC system. For reasons of maintainability, availability and uniformity, the stand-alone systems (which are required in 1993) will be constructed from commercial components that can emulate actual data acquisition protocols in software, with the few time-critical signals (especially trigger signals) generated by trigger modules supplied by the Trigger System group. It is hoped that many of the software standards that will be used in the eventual Data Acquisition and Online Computing Systems will be chosen in sufficient time such that so that the stand-alone data acquisition systems can have as much of the look and feel of the real Data Acquisition System system as possible.

Several institutions have begun investigating software tools and hardware workstations, interface and other components to be used in stand-alone data acquisition systems.

8.9.9. Quantities and locations of data acquisition system components

The following two subsections give tables of Data Acquisition System component quantities and locations for these components.

Component quantities

Table 8-11 lists Data Acquisition System prototype, pre-production and production component quantities. In addition to quantities in the table, miscellaneous test modules and designs, timing diagrams and behavioral models for the readout crate bus slave interface and event data readout port interface will also be provided by the Data Acquisition group to front-end and trigger system implementers.

Table 8-11
DAQ system component quantities.

Component	Prototype & Pre-production Quantity	Production Quantity	Total Quantity
SDC Standard DAQ Crates, Supplies & Cooling	26	288	314
SDC Standard DAQ CPU Modules	30	288	318
Crate Adapter/Interface Modules	30	288	318
Event Data Links	4	410	414
Event Builder Subsystem	1*	1	1* + 1
Event Data Flow Control Subsystem	2	2	4
Event Data Flow To/From Event Builder Message Links	2	2	4
Event Builder To Online Processor Data Links	8	36	44
Online Processor Subsystem	2*	1	2* + 1
Event Request Message Links	4	9	13
Online Processor To Online Storage Data Links	2	5	7
DAQ Control/Monitoring Network Routers	2	8	10
DAQ Control/Monitoring Network Repeaters	2	8	10
DAQ Control/Monitoring Network Links	10	327	337
Remote Control/Monitoring Links	9	136	145
Remote Control/Monitoring Link Repeaters	3	34	37
Electronics Rack Protection Subsystem	4	133	137

* Partial prototype

Physical Locations Of Components

Table 8-12 lists Data Acquisition System component locations. The Central Calorimeter, Outer Tracking Straw Tube and Muon front-end system readout crates are located within the body of the SDC detector, and are inaccessible during beam times. The rest of the readout crates are located in an upstairs area called the Counting Room.

Table 8-12
DAQ system component locations.

Component	Location
Inner Tracking Silicon Strip SDC Standard DAQ readout crates	Counting Room
Inner Tracking Silicon Pixel SDC Standard DAQ readout crates	Counting Room
Intermediate Tracking Gas Microstrip SDC Standard DAQ readout crates	Counting Room
Central Calorimeter (SCA) SDC Standard DAQ readout crates	On Detector
Central Calorimeter Shower Max SDC Standard DAQ readout crates	On Detector
Outer Tracking Straw Tube SDC Standard DAQ readout crates	On Detector
Forward Calorimeter SDC Standard DAQ readout crates	Counting Room
Muon SDC Standard DAQ readout crates	On Detector
Level 1 SDC Standard DAQ readout crates	Counting Room
Level 2 SDC Standard DAQ readout crates	Counting Room
Event Builder Subsystem	Counting Room
Event Data Flow Control Subsystem	Counting Room
Online Processor Subsystem	Counting Room

8.9.10. Commissioning The Data Acquisition System (Installation, Integration and Testing)

Installation of the first components of the SDC Data Acquisition System will occur near the end of 1996 (see Table 8-13). Continued installation, along with integration and testing of data acquisition components with those of other systems, will occur for a two year period and will be completed in early 1999. The Data Acquisition System allows for complete independence between front-end and trigger systems during commissioning.

The commissioning of the Data Acquisition System will occur in three phases. In the first phase, only the front-end and trigger system readout crates and the DAQ Control/Monitoring Network will be available. All front-end and trigger system components (cables, modules, etc.) will be tested using the DAQ Control/Monitoring Network connected to the DAQ CPUs in the readout crates. In the second stage, front-end systems will be tested at full data acquisition speeds, independently and separately from each other. This will be accomplished by either logically or physically dividing the Event Builder Subsystem and the Online Processor Subsystem. In the third phase, the entire Data Acquisition System will be tested for the specified functionality and performance.

8.9.11. System simulations

It is intended that all components of the Data Acquisition System be simulated at the system and subsystem levels using industry standard simulation tools such as Verilog and Modsim to gain a thorough understanding of both the buffering required to hold event fragments at various stages of the system and event data flow, its control and its throughput through the system.

The first important information to obtain from Monte Carlo simulations of the experiment is the event size and its distribution through the front-end systems and front-end readout crates. The maximum size of an event fragment fixes the buffer space (per trigger) required in the DAQ CPU and the data collection modules. The optimum number of buffers allocated in these readout modules depends on the data flow through the system.

The large spread in the event fragment sizes from front-end and trigger system sources (from a few bytes to a few kilobytes) requires data balancing and queueing at the input to the Switching Network of

the Event Builder Subsystem. Efficient implementation of this logic and the buffer space required for data balancing will be understood by simulating the Data Acquisition System architecture.

The Event Builder Subsystem and the Event Data Flow Control Subsystem will be modeled and simulated for a better understanding of the event data flow and its control. Both synchronous and asynchronous Event Builder Subsystem Switching Network modes of operation will be simulated.

The CPU time required to process events through the Online Processor Subsystem and the event queueing process within the Online Processor Subsystem are also subjects for simulations.

Current R&D activities are concentrating on modeling and simulation of front-end readout mechanisms and various Data Acquisition System architectures. These activities have been and are being conducted at the SSC Laboratory [67,68], Brown University, Fermilab [69], the University of Illinois [70] and the University of Toronto. The University of Illinois has also done additional work developing general simulation tools and standards for use by any group doing simulations. Further simulation studies are necessary to determine the optimum levels of redundancy needed in the DAQ architecture and the required reliability of DAQ components to insure that the data integrity and the quality of the physics will not be compromised.

8.9.12. Milestones and critical path items

As mentioned above, installation of the SDC Data Acquisition System will commence in 1996 and be completed in early 1999. The Data Acquisition System allows for complete independence between front-end and trigger systems during commissioning. During initial commissioning, only the front-end and trigger system readout crates and the DAQ Control/Monitoring Network will be available. Components such as the Event Builder Subsystem, the Event Data Flow Control Subsystem and the Online Processor Subsystem will begin to appear in mid-1997. Interfaces to the Online Computing System and the Control Computing System will also become available late in 1997.

The schedule for the data acquisition system is driven by the following needs:

- Stand-alone data acquisition systems must be available for test beam and lab bench runs starting in 1993.
- Initial data acquisition systems (complete with all readout and control components for limited numbers of channels) must be available starting in January 1997 for testing of front-end and trigger systems as the different detector systems are installed in the collision hall.
- The complete Data Acquisition System must be installed by January, 1999 to allow sufficient time for system integration and testing before data taking commences in October, 1999.
- Expensive components should be purchased as late as possible consistent with the above needs to minimize costs, as the cost of much of the data acquisition equipment is dropping with time.

These considerations result in the schedule of major milestones for the Data Acquisition System given in Table 8-13. Note that this schedule of milestones requires technology choices (for such items as the specific design of the standard readout crates and the Event Builder Subsystem) to be made during 1993.

The initial baseline design for the DAQ System has been described in the Data Acquisition System Conceptual Design Report [71]. This baseline design has been used for initial costing purposes and to demonstrate that there are sound technical solutions that meet the experimental requirements. It is expected that further R&D work over the next year will result in further improvements of the baseline design. Also, further work is needed to optimize the design where there are several different technology choices that could all meet the system requirements.

There are no major critical path items but it is important to determine early the protocols for readout of front-end electronics so design of the front-end integrated circuits, which do have a long lead time, can proceed. This may require some prototyping of DAQ System components to allow front-end systems to be tested. The only DAQ System items needing significant lead time are some of the components of the Event Builder Subsystem. The major development activity over the next two years consists primarily of software for Stand-Alone data acquisition systems for use in test beams and laboratories. Detailed component

requirements and specifications involving extensive discussions with industry also need to occur over the next two years.

Table 8-13
SDC data acquisition system major milestones.

Major milestones	Scheduled date
Completion of DAQ initial requirements document	November 1991
Completion of DAQ system design specifications, including technical choices	October 1993
Stand-Alone DAQ systems for use in test beams and labs	1993
Completion of DAQ component design	1994
Prototypes of all DAQ components	early 1996
Delivery of DAQ system for first front-end system	January 1997
Commissioning (installation, integration and testing) of complete DAQ system	Jan. 1997-Jan. 1999
Certification of a full, working DAQ system	July 1999

8.10. Ancillary systems controls

The Ancillary Systems Controls (ASC) system provides control and monitoring of subsystem and facility support equipment supplementary to that provided by the Control/Monitoring Network subsystem of the Data Acquisition system. In general, ASC maintains control of equipment that is located on the outside of the detector, in the detector hall and on the surface. The DAQ Control/Monitoring Network subsystem, in general, is concerned with the control, monitoring, parameter update and calibration of the electronics within the various Front-end subsystems and their communication with the DAQ and Trigger systems.

8.10.1. Architecture

The architecture of the ASC system is a multi-tiered structure of local computer systems interconnected by Local Area Networks (LANs). Each subsystem that is part of the total SDC instrumentation system has its own independent, autonomous local computer system. Typically, each is comprised of one or more control computers, a interface for a local operator, local communications and a set of I/O resources. Although independent, each computer system also communicates with an ASC supervisory computer system. Through the ASC supervisor, a supervisor operator can interact with each system in the same manner as can the local supervisor. External systems, such as Accelerator Controls or an online workstation, can also interact with the controls data system to monitor and, if granted permission, change operating parameters of the subsystems under ASC aegis.

The requirements of a local computer system may be met by a networked set of computers connected to a local supervisory control point or by a single computer acting as a control point to the local computer system's specific hardware. In either case, the minimum configuration for a local computer system is the local supervisory control computer. It provides a platform for control algorithm development, operator interface and the interrogation node for the ASC supervisory computers.

Implementations of the hardware generally follow the pattern of local control computer(s) and mass storage resources for each subsystem, and a graphics interface for development and operation and communications interfaces to the supervisory and lower level computers.

8.10.2. Software

The software systems developed by SSCL for use by the SSC accelerator will be also be used for the ASC and for the DAQ Control/Monitoring Network subsystem. It is intended to maximize the use of the software infrastructure being developed and supported by the SSCL and to have the SDC-supported efforts concentrate on the specific items related to subsystem ASC implementations, customized software and external interfaces.

SSCL core development

It is expected that the SSC software systems will provide at least the following features in its core software package:

- An interactive graphics interface for algorithm and control panel development.
- A distributed, LAN connected, shared-memory based data set for controls purposes.
- A modular design that facilitates customization for special requirements and the development of special-purpose interfaces.
- Uses industry standards for communications protocols and software languages.
- Support of hierarchical control and communications based on commercially available, industry-standard schemes.
- Capability to be implemented in configurations varying from stand-alone with no instrumentation to multi-tiered, hierarchical systems with diverse instrumentation support.

The ASC group will provide tools and periodic technical assistance and other support to the DAQ Control/Monitoring Network group, who will be adapting this software package for their applications. Their needs will include hardware specific algorithms, custom drivers and operator control screens utilized for their front-end electronics manipulations.

Instrumentation drivers

Special instrumentation driver routines will probably be necessary to cover the diverse needs of SDC ASC implementations. If the features of the SSC controls package are as expected, inclusion of a driver for a specific type of instrument should be straightforward. For example, a driver may be required for the interrogation of an vacuum ion gauge.

Many of the instruments available today have RS-232 ports with a special command set for control and read-back. Examples of current controls systems software—*e.g.*, TACL used at CEBAF—indicate that the driver routine could be relatively small, *i.e.*, ten's of lines of code—and could be easily integrated into the overall controls software. It is expected that the SSC controls core software system will support custom driver routines.

Interfaces to other systems

Communication between the ASC and other systems, such as the DAQ, the Online Computing System, and Accelerator Controls system, as well as users (experimenters) situated away from the SSC site, will be a vital need. The interfaces to all such systems will be implemented using a design with common features. These interfaces will be implemented on several levels.

The current design concept is to allow interrogation and set-point control by an external system through a dedicated LAN bridge containing bi-directional message interpretation code daemons. These daemon programs receive the LAN message traffic from the other system and are translated as required by the ASC software. The daemons then respond to the messages at a level of complexity dictated by the contents of the message—*e.g.*, set one power supply voltage, or set all power supply voltages to a predetermined array of set points. In general, the daemons communicate to ASC by accessing the distributed controls data set in the same manner as any other function of ASC does.

Access to ASC commands and data will be more or less restricted according to the needs of the external system, and safety and security requirements.

Online computing driver

It is expected that the Online Computing System will have unrestricted access, except for safety reasons, to any of the controls data set. Typical commands to the ASC system from online computing are; "Read the current gas flow through muon chamber XX" or "Turn off all high voltage on the central tracker." As in the case of the DAQ, the online computing interface will most likely reside in a dedicated LAN bridging computer. Again the specific functionality of the command interpreter will be defined in detail jointly by the ASC and Online Computing System groups.

External access driver

The purpose of the external access driver is to allow general access by authorized, pre-defined users. The external interface will allow them access to pre-defined commands and data sets, such as high-level summary information. Typically, they will have virtually no downward access to commands or data affecting the current operation of the detector.

SSC accelerator

The SSC accelerator interface requirements are a mix of the functionality of the DAQ or online interfaces with the restricted access features of the interface for external access. It is expected that the SSC accelerator controls will have access to any readable ASC parameter on demand. Its command subset will be very restricted, limited to commands related to overall safety or advisory information regarding the accelerator status. The details of the required commands will be defined in the future but the intent of ASC is to keep the allowable command set at a minimum.

8.10.3. Subsystem special requirements

Inevitably, each subsystem on the lowest architectural tier (Section 8.10.1) will have special needs. Some of these needs, together with descriptions of specific hardware required to accommodate them are discussed below.

Generally, ASC will provide the computer resources and instrumentation housing (i.e., crates) used in these applications. The specific instrumentation is provided by each subsystem in coordination with ASC.

The ASC group, working with the individual subsystems, will generate the control algorithms and operator control screens required to accomplish the control functions defined in the following paragraphs.

Silicon tracker

The Silicon Tracker has two sets of equipment to control and monitor the logic and depletion power supply, and the butane cooling system.

Both logic and depletion power supplies will be controlled and monitored by a single local computer communicating to 2 crates of instrumentation. The system will monitor all aspects of the power supplies operation, provide set point control, and maintain those set points through software algorithm feedback loops. Personnel safety for these supplies will be hard-wired into the power supply interlocks with the controls system monitoring and reporting the interlock status.

It is anticipated that the butane system, as acquired, will include its own set of controls and that the local ASC control computer will be able to communicate with it via a LAN connection. Software drivers will be developed as necessary. Control functions include monitoring and reporting the status of the system, controlling selected valves and issuing high-level commands such as "ON" or "OFF" or "PURGE." Additional functions may include passing anticipated heat-load information to the cooling system for feed-forward cooling algorithms.

Central tracker

The Straw Tube system will require monitoring and control of the high voltage power supplies and of the gas system. ADC's, DAC's, contact monitors and digital output modules supplied by the Straw Tracker group will be used. Personnel and equipment safety interlocks will also be monitored and reported by the controls.

The Fiber Tracker, if selected for SDC, requires the control and monitoring of the liquid helium system for the Visible Light Photon Counters. Temperature monitors, flow indicators, valve actuators, ADC's, DAC's, contact monitors and digital output modules supplied by the Fiber group will be used. Personnel and equipment safety interlocks will also be monitored and reported by the controls.

In either case, ASC will provide the computers, software and instrumentation crates necessary for the defined control functions.

Calorimeter and shower maximum detector

For the calorimeter and shower maximum detectors, ASC will provide controls for the phototube high-voltage power supplies and will monitor personnel and equipment interlocks. Two instrumentation crates connected to a local supervisor computer will perform the control functions required by this subsystem. Required instrumentation will be provided by the calorimeter and shower maximum subsystem groups.

Muon chambers and counters

The primary function of the muon chamber controls will be to monitor and control the high-voltage power supplies and the gas system for the chambers.

The subsystem will consist of eight local computers each with three instrumentation crates for monitoring and control of the high voltage system. An additional local controller is used to control and monitor the gas supply and mixture system. Personnel and equipment safety interlocks are monitored and reported.

Muon toroids

The controls for the muon toroids will monitor and control three major power supplies delivering approximately 3 MW to the barrel, intermediate and forward muon toroidal magnets. All instrumentation interlocks will be monitored and reported by the controls. One instrumentation crate will be required.

Superconducting solenoid

The superconducting solenoid has two items requiring control, the power supply and quench detection/protection system, and the cryogenic refrigerator plant.

One local computer and two instrumentation crates will be used to control and monitor the power supply and quench systems. All instrumentation will be provided by the superconducting solenoid group in consultation with the ASC group. The functions provided will include: precise control of the output current of the power supply; monitoring voltage taps, temperature and pressure sensors for quench calculations and monitoring and reporting of personnel and equipment safety interlock status.

The vendor of the cryogenic plant will supply a control system for the plant. Control and monitoring of the vendor's control system will be executed by the ASC control computer. As in the case of the butane cooling system of the Silicon Tracker, the required operations are mainly establishing setpoints and monitoring of the cryogenic plant status.

Data acquisition

Two levels of interfacing will be provided between the ASC and DAQ systems. One level is based on a LAN interrogation/command protocol, the other on hard-wired permissives.

The LAN interface consists of a dedicated processor acting as an active command interpreter and gateway. It is faceless and has at least two separate LAN ports, one connected to the ASC supervisory LAN and the other connected to the DAQ command LAN. It is anticipated that this computer will be a diskless node supported by one of the computers in the ASC supervisory console system.

The other interfacing unit provides digital hard-wired signalling between the DAQ and ASC. Typically, this will be used for faster notification of conditions within a predetermined set—*e.g.*, personnel interlock violations, power supplies out of tolerance etc. The complete set of conditions is yet to be finalized but will be represented by no more than 50 bidirectional signals. One computer with an operator display and one instrumentation crate will be required. The function of this system is to monitor across the entire ASC, picking up relevant information and calculating the permissive formulas. Each permissive calculation results in an output to the DAQ. Instrumentation resources for this function is provided by ASC.

Online computing

The interface to the Online Computing System is a faceless computer with at least two LAN ports, one connected to the ASC supervisory LAN and the other to the online LAN. The function of this interface is to interpret commands from the online systems to access the ASC data set. As in the case of the DAQ LAN interface, this unit will be a diskless node supported by one of the supervisory console computers.

Supervisor

The function of the Supervisor system is to provide access to the central data set, integrated control across the subsystems, and the operator interface. It consists of four computers, each with at least one large screen, color monitor and appropriate mass storage systems. Each computer has two LAN ports, one connected to the ASC supervisory LAN and one connected to the External Access LAN.

Environmental monitor

The environmental monitor maintains surveillance over critical systems and spaces of the SDC. Examples include: monitoring of front-end crate temperatures; water flow; voltages and currents; and monitoring temperatures internal to the detector that are of significance to more than one subsystem.

This system consists of one computer with operator display connected to as many as twenty instrumentation crates that are remotely located throughout the detector and hall.

Utilities monitor

The basic function of the Utilities Monitor is to provide an interface between the SDC installation and the rest of the SSC facility. The function of this monitor is not well defined as yet, but will include transfer of global information such as: power monitors, fire alarms, HVAC monitoring etc.

The monitor will consist of one computer with an operator interface connected to twenty remotely located instrumentation crates.

Safety monitor

The safety monitor provides oversight of personnel access, Oxygen Deficiency Hazard alarms, radiation monitors located throughout the SDC installation. It reports to the other subsystems controls and can issue commands to mitigate a perceived hazard.

The monitor consists of a computer with operator console connected to twenty remotely located instrumentation crates.

Developmental system

The developmental system provides an offline environment for the development of algorithms, custom programs, operator control screens etc. It will consist of six workstations interconnected on a LAN together with resources for mass storage, archiving and printing. Three of the workstations each have three instrumentation crates attached.

The intent of this system is to provide a means to develop applications software for various subsystems prior to the arrival of their hardware. It also provides equipment to test instrumentation and control concepts. It has enough hardware depth that simulations of actual in-field performance can be accommodated. A selected amount of generic instrumentation will be provided by ASC. Other instrumentation resources will be provided by the subsystems.

Core program maintenance, programming assistance and general consultation support services will be provided as part of this system.

References:

1. A. Bhattacharya, *et al.*, "Front End Electronics Development for SSC Detectors," (unpublished) (October 2, 1989).
2. A. Bhattacharya, *et al.*, "Front End Electronics Development for SSC Detectors," (unpublished) (October 1, 1990).
3. C. Blocker, *et al.*, "Front End Electronics Development for SSC Detectors," UPR-208E (October 1, 1991).
4. "Silicon Tracker Conceptual Design Report," SCIPP 92/04 (Jan. 1992).
5. "Conceptual Design: Intermediate Angle Track Detector," SDC-92-171 (Jan. 1992).
6. "Conceptual Design Report for the Modular Straw Outer Tracking System," SDC-91-00125 (Jan. 1992).
7. "Conceptual Design: Scintillating Fiber Outer Tracker," SDC-92-174 (Jan. 1992).
8. "Calorimeter Electronics Preliminary Conceptual Design and Progress Report," (unpublished) (Aug. 1991).
9. "Baseline Implementation of the Digital Phototube Readout System for the SDC Calorimeter," (unpublished) (Jan. 1992).
10. John Oliver, "Muon System Front End Electronics Preliminary Conceptual Design Report," (unpublished) (Oct. 1991).
11. A. Lankford, *et al.*, "Conceptual Design of the SDC Trigger from the SDC Trigger Group," SDC Note SDC-91-098 (1991).
12. T. Gorski, J. Lackey and W.H. Smith, "The First Four Microseconds: SDC Level 1 Pipeline Length," SDC Note SDC-91-139 (1991).
13. W.H. Smith, T. Gorski and J. Lackey, "SDC Trigger Preliminary Conceptual Design," SDC Note SDC-91-089 (1991).
14. G. Sullivan and M. Miller, "Calorimetry Trigger Rates at the SDC" SDC-91-099 (1991).

15. T. Gorski, J. Lackey, W.H. Smith and W. Temple, "Isolated Electron Pattern Logic Design and Performance at the SSC," SDC Note SDC-91-087 (1991).
16. Zeus Collaboration, "Status Report 1989," DESY, Hamburg, Germany (1989). Fig. 5.8.10.
17. T. Gorski, J. Lackey and W.H. Smith, "Level 1 Trigger Decision Design for the SDC," SDC Note SDC-91-088 (1991).
18. T. Gorski, J. Lackey and W.H. Smith, "Revised Level 1 Trigger Decision Design for the SDC," SDC Note SDC-92-198 (1992).
19. G. Sullivan, "A Proposal for the SDC Level 1 Trigger Design," SDC Note SDC-91-100 (1991).
20. T. Gorski, J. Lackey and W. H. Smith, "Correlated Level 1 Muon Trigger Design for the SDC," SDC Note SDC-91-141 (1991).
21. M. Thompson, "Fiber Optic Data Transmission for the SDC Detector," SDC Note SDC-91-092 (1991).
22. D. Panescu, T. Gorski, J. Lackey and W.H. Smith, M. Thompson, "A Pipelined CMOS 4x12-bit Adder," SDC Note SDC-92-218 (1992).
23. W. Temple, T. Gorski, J. Lackey and W.H. Smith, "SDC Jet Trigger Studies," SDC Note SDC-91-140 (1991).
24. W.H. Smith, T. Gorski and J. Lackey, "SDC Global Level 1 Processor: Clock & Control," SDC Note SDC-91-090 (1991).
25. S.R.Amendolia, *et al.*, "Study of a Fast Trigger System on Beauty Events at Fixed Target and Colliders," Nucl. Instr. and Meth. **A289**, 539 (1990).
26. B.;Denby, *et al.*, "Neural Networks for Triggering," IEEE Trans. on Nucl. Sci. NS **37**, 248 (1990).
27. R.M.Lea, "ASP: A Cost Effective Parallel Microcomputer," IEEE Micro, (Oct. 1988).
28. R.K.Bock, *et al.*, "Feature Extraction in Future Detectors," Nucl. Instr. and Meth. **A289**, 534 (1990).
29. J-C. Brisson, *et al.*, "The SDC Level 2 trigger," (SDC note in preparation) (April 1992).
30. C. Bebek, *et al.*, "Scintillating Tile Barrel and Endcap Calorimeter and Shower-Maximum Electronic Cost Estimate," (unpublished) (Sept. 4, 1991).
31. C. Bebek, *et al.*, "Front End Electronics for SSC Calorimetry Progress Report," (unpublished) (Aug. 16, 1991).
32. S. Kleinfelder, IEEE Trans. Nucl. Sci. NS-**37** (1990) p. 1230.
33. S. Kleinfelder, M. Levi and O. Milgrome, "Test Results of a 90 MHz Integrated Circuit Sixteen Channel Analog Pipeline For SSC Detector Calorimetry," LBL-30222.
34. S. Kleinfelder, M. Levi, and O. Milgrome, "Test Results of a 90 MHz Integrated Circuit Sixteen Channel Analog Pipeline For SSC Detector Calorimetry," LBL-30222.
35. S. Kleinfelder, M. Levi and O. Milgrome, "Toward a 60 MHz Analog Virtual Pipeline Integrated Data Acquisition System, Nucl. Phys. **B23A**, 382 (1991).
36. O. Milgrome, S. Kleinfelder, and M. Levi, "A 12 Bit Analog to Digital Converter for VLSI Applications in Nuclear Science," (submitted to IEEE Trans. Nucl. Sci.).
37. "Calorimeter Conceptual Design: Tile/Fiber Scintillator Option," SDC Tile/fiber Calorimeter Group (Sept 3, 1991).
38. V. Barnes, Ft. Worth SSC Symposium, (1990).
39. J. Hoff, "A Full Custom, High Speed Floating Point Adder," FNAL report (Jan 10, 1992).
40. M. Larwill, *et al.*, "Test Results on the Floating Point Adder Test Chip," FNAL report (Jan 1992; updated Feb 1992).
41. T. Zimmerman, "PMT Digitizer Current Splitter Test Results," FNAL technical note (Jan 14, 1992).

42. Analog Devices Inc., Norwood, MA, #AD9020, #AD9040 and similar devices in development; also Signal Processing Technologies Inc. #SPT7814 and related devices in development.
43. L. Mo, *et al.*, "Cockcroft-Walton Phototube HV Supply," VPI-IHEP-91/1 (June 1991)
44. M. Haldeman and R. Yarema, "Test of Digital Transmission Characteristics of 0.025" Pitch Flat Cables," (Jan 23, 1989).
45. CDF Collaboration, Nucl. Instr. and Meth. **A267-271** (1988).
46. B. Baumbaugh, *et al.*, "Studies of Visible Light Photon Counters with Fast Preamplifiers", (submitted to the IEEE Trans. Nucl. Sci.)
47. T. Zimmerman, IEEE Trans. Nucl. Sci. **NS-37**, No.2, 439 (1990).
48. R. Yarema, *et al.*, "A High Performance Multi-Channel Preamplifier ASIC," (submitted to the IEEE Trans. Nucl. Sci.)
49. F.M. Newcomer, R. Van Berg, J. Van Der Spiegel and H.H. Williams, "High Speed Bipolar Integrated Circuits for SSC Applications," Nucl. Instr. and Meth. **A283**, 806 (1989).
50. R. Krummenacher, C. Enz and R. Bellazzini, INFN PI/AE 91/08.
51. R. Partridge, *et al.*, "SDC Data Acquisition Design Requirements and Functional Specifications," SDC Data Acquisition Working Group document (September 1991).
52. Ancor Communications, Inc., Minnetonka, Minnesota.
53. Network Systems Corporation, Minneapolis, Minnesota.
54. UltraNet, San Jose, California.
55. "Fibre Channel Physical Layer," ANSI X3T9.3/91-071, Rev. 2.1, May 1991.
56. "High Performance Parallel Interface (HiPPI), Mechanical, Electrical and Signalling Specifications," ANSI X3.183-1991, Rev. 7.1 (February 1990).
57. "Serial-HiPPI," Serial HiPPI Implementer's Group, Rev. 1.0 (May 1991).
58. "Synchronous Optical Network (SONET) Transport System: Common Generic Criteria," Bellcore (Bell Communications Research), TR-TSY-000253, Issue #1 (September 1989).
59. "Synchronous Optical Network Add-Drop Multiplex Equipment (SONET ADM) Generic Criteria," Bellcore (Bell Communications Research), TR-TSY-000496, Issue #2, (September 1989).
60. E. Barsotti, *et al.*, "A Proposed Scalable Parallel Open Architecture Data Acquisition System for Low to High-Rate Experiments, Test Beams and All SSC Detectors," IEEE Trans. Nucl. Sci., **NS-37**, No. 3 (June 1990)
61. M. Bowden, "System Overview of the Scalable Parallel Open Architecture Data Acquisition System," (unpublished) (May 1990).
62. M. Bowden, *et al.*, "Scalable Parallel Open Architecture DAQ Prototype Project Report," (unpublished) (November 1991).
63. E. Barsotti, *et al.*, "A Progress Report Of The Switch-Based Data Acquisition System Prototype Project and The Application Of Switches From Industry To High Energy Physics Event Building," SSC R&D Conference, Fort Worth, Texas (October, 1990).
64. A. Booth, *et al.*, "Software Description Of The Scalable Parallel Open Architecture Data Acquisition System," internal Fermilab report, (August 1990).
65. "Scalable Coherent Interface (SCI)," IEEE P1596.
66. International Business Machines, Kingston, New York.
67. M. Botlo, *et al.*, "Data Collection Studies For The Superconducting Super Collider," Nucl. Instr. and Meth., Elsevier Science Publishers (February 1992).

68. E. C. Milner, *et al.*, "Data Acquisition Studies For The Superconducting Super Collider," IEEE Trans. Nucl. Sci., (to be published in June 1992).
69. A. Booth, *et al.*, "Simulation And Modeling Of Data Acquisition Systems For Future High Energy Physics Experiments," IEEE Trans. Nucl. Sci., (1990).
70. E. Hughes, *et al.*, "Modeling And Simulation Of The SDC Data Collection Chip," IEEE Trans. Nucl. Sci., (to be published in June 1992).
71. E. Barsotti, *et al.*, "SDC Data Acquisition System Conceptual Design," SDC Data Acquisition Working Group, (November 1991).
72. Silicon Tracking Conceptual Design Report (Jan. 1992).

9. Online computing

9.1. Overview

Chapter 8 described five parts of the SDC detector data collection system—front-end electronics, trigger system, DAQ/event builder, Level 3 trigger processors, and slow controls. The complex tasks of supporting, configuring, managing, and monitoring these systems so that reliable data recording is assured are the core responsibilities of the online computing system. Its primary goal is to record data on permanent storage media in a form suitable for offline reconstruction. The online system needs to control and record a sustained data rate of up to 100 Mbyte/s. It also provides services, resources, and a general task framework for the support of data-taking. A high degree of reliability must be achieved by extensive monitoring not only of the data, but of all relevant parameters and processes that affect data taking. If offline reconstruction requires further information than that recorded on permanent media, the online system is incomplete. If data are lost or corrupted in some way, then the online system has failed. While failures may be attributable to other systems—DAQ, ancillary controls, *etc.*—the online system as overall supervisor must deal with problems in an expeditious manner.

The online system consists of a number of workstations serving as dedicated computers, file and database servers, and as operator interfaces for control and monitoring. It also includes a large disk space for databases and data spooling, and the mass storage system for permanent data storage.

The online computing system is primarily a complex software system. We estimate that approximately 50 man years are required to produce the software system.

9.2. Functional requirements

9.2.1. Acquiring and recording data

Operability. Operation of the detector should require two physicists who do not have special skills in online operation.

Reliability. The system should contribute no more than 5% to the total deadtime. This does not include rate dependent dead-time introduced by DAQ and Trigger systems. From time to time, it will be necessary to re-configure the online system. This should take no more than five minutes. Restarting the system should take less than ten minutes and a cold start, requiring rebooting and loading of DAQ, trigger system and front-end parameters, should be accomplished in less than 30 minutes. We recognize that there may be an initial shakedown period where these requirements are not met.

Data logging. It will be necessary to record a maximum sustained data rate of 100 Mbyte/s out of the Level 3 processor farm. The data will be split into streams tagged by the Level 3 trigger event classification. One stream of interesting events will be directed to express offline reconstruction. Another type of data stream will contain all time varying experimental parameters and conditions, such as detector voltages and calibrations.

Run control and human interface. High level run control should be highly automatic. The human interface for run control should be location independent, although certain actions may be restricted to specified locations. It should be possible to interact using simple command lines, point-and-click, function keys, touch panels, *etc.*

9.2.2. Insuring data integrity

Data monitoring. The online system should monitor data emerging from Level 3. This should be implemented as a group of processes that receive specific kinds of events from Level 3 for monitoring.

Parameter and configuration management. Changes in the detector electronics and online system configuration may occur between runs. A complete specification of the configuration must be recorded on disk.

Calibration services. It should be possible to configure the detector for calibration and to initiate automatic sequences of calibration events. A means of viewing the history of calibration constants for a single channel or an entire sub-detector should be provided.

Error detection and alarms. Alarms can be handled at more than 10 Hz. Errors can be logged at more than 100 Hz.

Diagnostic services. The online system software framework should facilitate and support such diagnostic codes as may be required for debugging sub-systems and debugging the online system.

Database/history services. Online databases will include recent sets of configuration parameters, full run summaries for the past month, and brief run summaries for the entire life of the experiment.

9.2.3. Physics data processing

Interactive display and analysis of pools. Facilities will be provided for the interactive display and analysis of events following the Level 3 processors. According to physicist-defined criteria, data from the Level 3 output streams are replicated in sharable event pools.

Dynamic full analysis framework. A full off-line analysis framework should be provided for interactive analysis of reconstructed events. This framework should support dynamic insertion of user analysis routines and be capable of analyzing previously reconstructed data.

9.3. Implementation

Here we describe the implementation of the online system that will satisfy the requirements described above.

9.3.1. Methodology

The highest level components of the online system will be defined at an early stage by a top-down design using standard software engineering tools. A limited number of programming languages will be used; C++ and FORTRAN 90 probably will be used for the bulk of the online system source code. A great deal of hardware-specific, bottom-up software will proliferate as hardware is acquired. To merge the top-down design and the bottom-up implementations will require active participation by persons involved in the online system coding. Online system coding will make use of rapid-prototyping so that programmers can quickly develop familiarity with the hardware, recognize new requirements and eliminate impractical ones. This will also permit integration of the online software installation schedule with the detector installation schedule. Development and debugging of the hardware/software interfaces will proceed long before the software or hardware have passed their individual acceptance tests.

Many problems of operability and reliability occur at system interfaces. We will endeavor to identify these interfaces and to minimize their number and complexity, both between online system processes and their interactions with other systems. Neither hardware nor software can be 100% dependable. When problems occur and are understood, they will be either corrected in the code or an automatic corrective work-around implemented.

The groups developing the front-end electronics, trigger, and DAQ systems will require frequent consultation with the online developers to avoid bottle necks in down-loading configuration parameters and in booting many CPU's from a single server. This is to guarantee that the online system specifications for down time and warm and cold start times can be met.

Incremental builds of the online system will provide important checks on the progress of the integration of hardware-specific software. They allow periodic checks of the compatibility and scalability of the entire data collection system.

9.3.2. Functional implementation

Figure 9-1 shows the data paths within the online system. The nodes on the paths are processes or groups of processes running in the online system. They are not necessarily bound to a particular CPU. Figure 9-2 is a schematic of the control paths. The workstation consoles are the sources of human control as well as the primary providers of visual information about the data-taking process.

A process in the online system will have a standard outer shell which may consist of a data interface for receiving and sending large data segments (*e.g.*, events), a control interface for receiving and sending short textual control messages, and a graphical user interface for displaying data and receiving control input from a console. In general several processes will share a single CPU. The data interface performs operations such as queuing of buffers to each process within a processor as well as management of the flow of buffers into and out of the processor. User code is encapsulated in these template programs. Extensive use of a UNIX-like *make* facility will ensure that pieces of the online system are up-to-date with respect to one another.

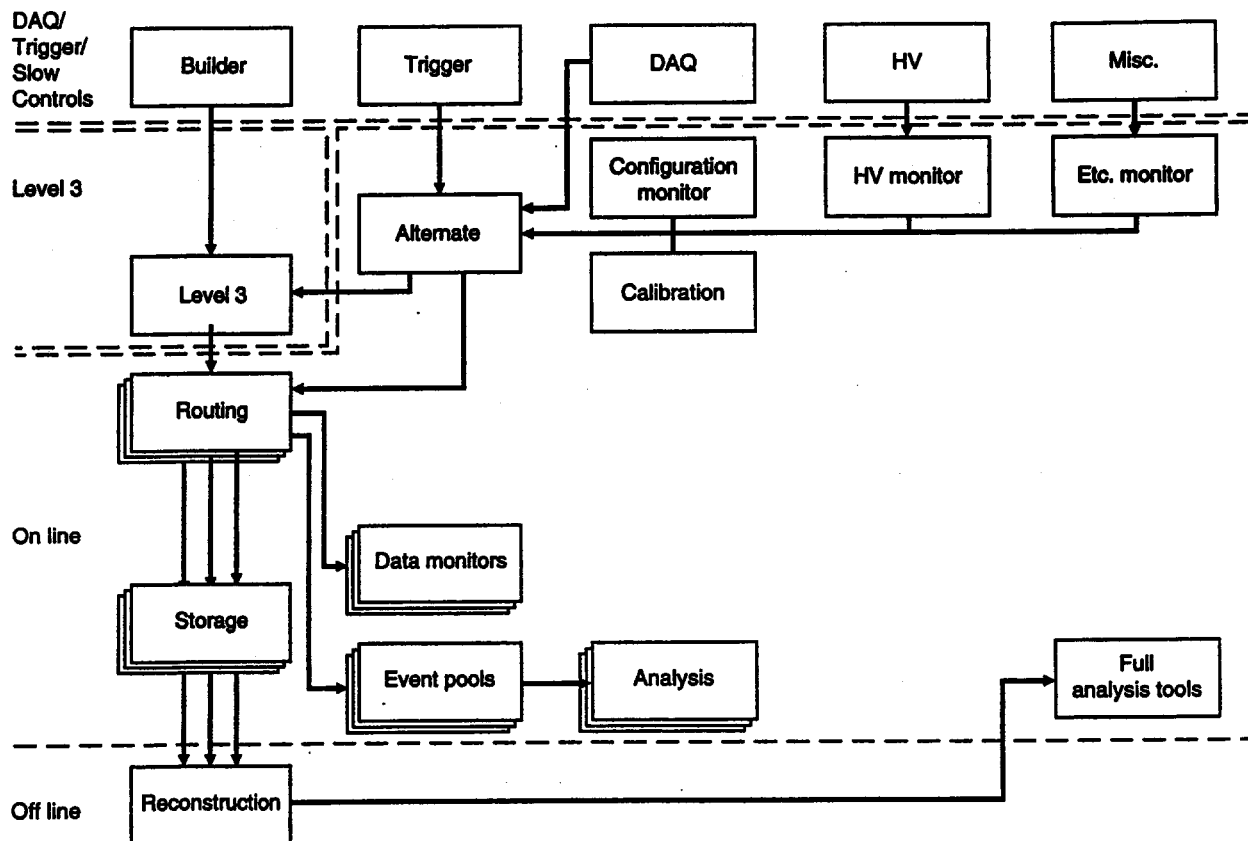


FIG. 9-1. Data path connections within the online system and with other detector systems.




FIG. 9-2. Control paths between online system components and other detector systems.

9.3.3. Hardware

Online system processors will be commercial workstations of about 100 MIPS processing power configured with 4 Gbyte of disk, 100 Mbyte memory, a console, and various network and serial ports. Higher performance workstations may be required in certain parts of the online system.

9.3.4. Run control

High level run control is the responsibility of the online system. This will be highly automated, perhaps employing an expert system to decide actions and responses. The human interface for run control will be location independent, although certain actions might be permitted only from specified locations. It may be possible for domestic and foreign collaborators to have mini-control-rooms at their home institutions so that they can monitor detector performance and staff data taking shifts. We are actively investigating the feasibility of a remote control center in Japan that would be linked by high speed data and video networks to the local control room at the IR site. Although the local control room would always be staffed during detector operation, the ability to take shifts from remote sites is an attractive possibility. If our study demonstrates the feasibility of the concept in Japan, such remote centers might be replicated elsewhere in the U.S., including the SSCL central campus, and in other countries.

A high level screen management package will be used so that modifications to the human interface will be mostly confined to simple edits using a screen/display editor. An online system programmer will be able to implement a new screen display within a few hours. Screens will perform the function of data display and operator input device. The display package will be implemented with non-proprietary software so that it can be broadly distributed within SDC. This will enable collaborators to monitor and interact with the online system from their home institutions.

Additional information will be displayed on TV monitors from remote cameras and from computer generated realtime images. The online system will be capable of generating audio messages as well as text.

The primary executor of run control is the XPC process (see Fig. 9-2). XPC receives status messages and commands from all processes in the online system as well as external systems. In particular, it receives control messages from the operators' console windows. It may generate its own control messages in response to changing experimental conditions and may communicate with an expert system to determine appropriate responses.

9.3.5. Configuration control and monitor

Configuration control and monitor software is responsible for providing all necessary down-loadable parameters to the DAQ and trigger systems. It also is responsible for configuring the online system for data-taking, calibration, debugging, *etc.* The configuration and performance of the online system will be continuously monitored. When exceptions are detected it is the responsibility of this software to notify XPC and to take corrective action.

A complete description of the present configuration must be maintained on disk. Previous configurations will be preserved in a database. The size of a description is dominated by front-end electronics parameters—pedestals, gains, *etc.*—which may consist of 10 data words per detector channel. Each configuration will occupy approximately 400 Mbyte of disk space.

9.3.6. Data logging

The online system will be capable of recording up to 100 Mbyte/s of event data—100 Hz of 1 Mbyte event data output by Level 3. Event sizes up to 20 Mbyte will be supported to handle the large records that can result during calibration.

After Level 3 trigger processing, the surviving events are routed to appropriate streams and data monitoring processes according to the classification of the event by Level 3. There will be ~20 streams for recording on permanent media. The streams may not be mutually exclusive at the 10% level. An express stream will accept data at up to 1 Hz for immediate full reconstruction.

All other parameters necessary for offline reconstruction will be recorded as necessary in a separate stream. These are the data for experimental conditions—detector HV's, magnet current, temperatures and environmental, configuration parameters, *etc.* We estimate that this stream will be less than 100 kbyte/s.

It will be necessary for the online system to record events from sources other than Level 3. For example, during debugging of a front-end system, it may be desirable for the crate resident CPU to record data on an 8 mm tape drive. An alternate event source will be provided to manage such data, to inject data into Level 3 for testing, and other such tests as may be necessary.

Mass storage system

Data from Level 3 is buffered onto a spooling disk before being copied to the permanent storage system. Offline processors sometimes require immediate access to the data and this can be accomplished by reading the spool disks. These disks are connected to a few workstations or may be attached directly to the online system network. Disk striping and RAID technology will permit disk transfer rates sufficient to utilize 30 Mbyte/s tape drives effectively. We estimate the required spooling space to be 100 Gbyte.

The mass storage system for the online system probably will be a robotic tape storage system capable of recording 30 Mbyte/s on each drive. Such storage is available today and we anticipate that cost and performance will be more favorable in the future. This system can be inexpensively remoted from the IR site if fiber optic networks are used.

Data routing subsystem

There are approximately 1000 Level 3 processors. This requires a high degree of connectivity to the data spooling disks and the permanent storage. This can be accomplished by several means which fall into two basic classes: switching or network routing. High speed switches of sufficient capacity will be available at reasonable cost by 1995. SONET technology, which will be used in the SSC accelerator control network, supports 2.5 Gbit/s data rates using optical fiber over quite large distances (up to 40 km). The event builder subsystem within the DAQ is considering Fibre Channel (1 Gbit/s) to deliver the data to the Level 3 processors and it may be desirable to continue this network into the online system. Industry developments are being followed closely.

9.3.7. Data monitoring

The integrity and contents of the data will be monitored as it is acquired and written to the mass storage system. The online system must be capable of routing specific kinds of data to different monitoring processes. These processes will be capable of looking at about 1 Hz of events and the number of such processes will correspond to the number of sub-systems requiring such monitoring.

9.3.8. Condition monitoring

Some additional processes are required to monitor non-event data—voltages and environmental conditions—as well as the operation of the online system itself. Some of these will operate at rates of up to 50 Hz. These processes will operate even in the absence of colliding beams since they may be parts of feedback systems. Failsafe processes will operate locally within the slow control systems.

Interlocks and failsafe systems are in general outside the domain of the online system. However, the online system will monitor the status of any interlocks that are important for data-taking. Additionally, the online system will redundantly interlock certain high voltages with accelerator beam status and retain high level control of the HV system. Other safety and failsafe systems will be implemented independently as part of the SDC slow-controls or accelerator control systems.

There will be a standard error logging and detection service which can handle 10 Hz of alarm messages and 100 Hz of logged error messages. An alarms-and-limits package will provide a standard method of detecting abnormal conditions through comparisons of current and historical conditions. This package will be capable of tracking the value of a variable over time and determining significant variation in a general way, such as by comparison of a histogram to a standard histogram.

9.3.9. Calibration

Software will be provided to manage calibration data and calibration sequences. Sometimes it will be necessary to log calibration data for offline processing and at other times calibrations will be processed in the online or DAQ systems and only the derived constants logged to the offline system.

9.3.10. Physics data processing

Users will have the capability of dynamically defining selection criteria to create event pools of events which can then be shared by several users who wish to do physics analysis. The online system will be capable of supporting up to 10 event-pool streams with an individual rate of no more than 0.5 Hz and an integrated rate of no more than 2 Hz. There will be five workstations provided for interactive analysis, viewing of event-pool data, and graphical display of events.

A single stream of fully reconstructed events from offline processing or possibly from Level 3 processing will be provided to use physics processes to monitor the detector performance. This stream will provide events up to a maximum rate of 1 Hz.

9.3.11. Database

The online system will require a database to store the details of recent configurations, full run summaries for recent runs (over the past month), and brief run summaries for the duration of the experiment. We anticipate that this will require 100 Gbyte of disk space. The alarms and limits package may utilize this database.

9.3.12. Documentation

Effort will be required to document the online system code and operational procedures. The human interface will provide prompting and help when the user requests it. It will be difficult to maintain up-to-date documentation in a rapidly evolving system. We will use self-documenting code so that some documentation can be generated automatically. Documents will be managed using the code management system.

9.3.13. Software tools and standards

The online system will require standard software packages for code management, data management, code development, data bases, graphics, and human interfaces. Where appropriate, the same standards will be found in the offline system.

9.4. Relationships and impact on other systems

The online system interacts with a number of detector sub-systems. Generally, it acts as a host which provides services to the subsystems: cross-compilers, configuration parameters, file management, and archival storage. The online system is in direct control of the DAQ, trigger hardware, ancillary controls, and Level 3 processor farm. It communicates with or contains user code from all of the various detector sub-sections. Contact with offline computing is through the mass storage system.

Care must be taken to insure that compatible interfaces exist between these sub-systems. Points of contact will be minimal and simple so that parallel development may proceed in each of the sub-systems without excessive interdependence requiring an unrealistic degree of cooperation and organization. On the other hand, it will sometimes be advantageous to share standards across systems so that effort will not be duplicated. This philosophy extends even inside the online system software where interdependence of processes and code will be minimized. Cases where effort can be reduced are often identified naturally as development proceeds. The points of contact must be specified at an early date so that development will not be held up and effort wasted. We briefly summarize the relationships with other systems below.

9.4.1. Front-end electronics

The online system will archive relevant parameters and be capable of supplying them to the front-end electronics on demand. It will be possible to login to processors in the front-end crates from online computers. Coding support for these computers resides in the the online system.

9.4.2. Level 1 and 2 trigger systems

The online system will archive relevant parameters and be capable of supplying them to the trigger systems on demand. It will be possible to login to processors in the trigger systems. A trigger monitor process will run in the online system containing code supplied by the trigger group. The online system is also responsible for high level trigger control and configuration. The online system can receive and log data directly from the trigger system as necessary.

9.4.3. Data acquisition and Level 3

The online system will archive relevant parameters and be capable of supplying them to the DAQ systems on demand. It will be possible to login to processors in the DAQ systems. A DAQ monitor process will run in the online system containing code supplied by the DAQ group. The online system is also responsible for high level DAQ control and configuration. The online system can receive and log data directly from the DAQ system as necessary. The online system will have one process dedicated to high level control of the Level 3 processing farm. Routing of data after Level 3 processing may require the online system to communicate routing parameters to Level 3 processors.

9.4.4. Offline computing

The online system must provide complete data on permanent media for offline analysis. It is the responsibility of the online system to insure that the data is complete and uncorrupted. Occasionally, problems will be detected offline which require online correction. Online programs will be capable of receiving messages from offline programs. The offline system will also be capable of providing a stream of reconstructed events for online monitoring at rate of up to 1 Hz. It will be possible to login to offline processors from the online system.

9.4.5. Test beam systems

Some parts of the online system will exist as prototypes used in test beam DAQ systems. It will be natural to import some user codes developed during test beam running and adapt them for colliding beam running.

9.4.6. Ancillary controls

The online system will archive relevant parameters and be capable of supplying them to ancillary control systems on demand. It will be possible to login to processors in the slow control systems from online computers. The online system will have a few condition monitoring processes which will primarily access data from slow controls concerning various environmental parameters. The online system will be able to both monitor and control HV's in the slow control system.

9.4.7. User code

The online system must be able to run, control, and monitor various types of user code. These codes appear in the data monitoring processes, condition monitoring processes, configuration monitoring and control processes and other sub system control processes. The online system will provide standard templates to encapsulate user code so that externally it has an online appearance and functionality.

9.5. Schedule

The schedule for prototyping and testing of the online system is closely linked to the DAQ installation schedule. Some major milestones are the next test beam running at Fermilab in 1994, the beginning of the test beam program at the SSCL in January 1997, first installation of muon chambers completed in May 1997, cosmic ray running of the detector in March 1999, and first data-taking in October 1999.

9.5.1. Test beam testing

Some parts of the online system software will be tested as prototypes in the SDC test beam DAQ systems. Initial tests will be made in 1994. By 1996, tests will include the prototypes of the data-interface, the control message interface, the online process template, human interface, and alternate data source systems. A prototype experiment control process (XPC) will also be needed in the test beam system. These prototypes will be substantially simpler than the versions required for colliding beams, but will contain the kernels of the necessary software and serve as checks of the basic design. Prototypes of the histogramming and alarms and limits and HV packages may also be tested. Development of all these

prototypes is estimated to take seven man years. First versions would be available for Fermilab running in 1994 and an SDC version would be released in early 1996. This would provide an early hardware/software integration test.

9.5.2. Data format

The format of data used in Level 3 processing and subsequently logged to tape will be specified at an early date. The specification of this format will permit offline simulations to produce realistic data for testing of the online system and permit timely development of online software which needs to have this defined. We require that this format be specified by mid-1993.

9.5.3. Sub-detector installation

The first muon chamber installation will be complete in May 1997. DAQ and online groups will implement the readout of these chambers. The online system will experience an important incremental build and test at this time. The system will not as yet need an operational Level 3. Nor will it need a fully developed data-routing network. The alternate data source will be able to handle testing in the sub-detector installation phase. After mid-1997, installation of sub-detectors occurs every few months requiring frequent additions to the growing DAQ and online systems.

9.5.4. Level 3, data routing and logging

Prior to the cosmic ray running scheduled for March 1999, complete Level 3, data routing and logging must be operational. Delayed purchase of the hardware for these items will result in significant cost savings. However, prototypes of Level 3 and the data routing system will exist by late 1997 so that online software can be developed and tested.

9.5.5. Cosmic-ray running March 1999

Cosmic ray running of the entire detector after March 1999 is the first shakedown of the complete online system. It will test most parts of the online system. Final adjustments will follow in order to be prepared for real data taking in October 1999.

9.5.6. Detector turn-on October 1999

During accelerator commissioning, XPC logic for automatic running will be developed and tested. Early data taking will probably lead to substantial fine-tuning of such logic before truly automatic operation can be achieved. We are likely to recognize some new requirements at this stage as the entire detection system responds to real rates and quantities of data. The online system will be modified as practical to meet these requirements. Extensive tuning of user code dealing with calibration and fault detection will be expected. Tuning of the data routing network for real data will also be required.

9.6. Organization

9.6.1. Project management

Online system software will be written by a small, coherent group, centered at SSCL. A larger group of physicists and computer professionals from throughout the collaboration will participate in specification and design reviews to insure a broad perspective. However, the actual coding is expected to be done by a dedicated online group at SSCL.

Prototype versions of online code will be evaluated during various test-beam activities in the years prior to experiment turn-on, allowing many members of the collaboration to gain experience with and comment on the system.

Various online tools and application packages will be written by a larger collaboration-wide online group. These tools and packages will be tested and integrated by the dedicated SSCL online group. This

group will also provide support for the testing and integration of application packages written by subsystem physicists to run in the online environment. Final responsibility for these dedicated subsystem applications will remain with the subsystem physicists.

10. Offline computing

10.1. Introduction

The SDC offline computing system will provide the capabilities needed by physicists for production, analysis, and simulation; developers for software design, coding, testing, integration, and maintenance; and system administrators for management and control of the system configuration, change/update management, problem response and system recovery. Given the data rates, event size, and number of physicists expected in the SDC, the system must handle raw data rates and supply processing capability, storage capacity, and data access that are two to three orders of magnitude greater than corresponding systems for detectors now in operation. The size and geographical distribution of the SDC require extensive network connectivity and high speed data transfer capability. These capabilities can be provided by a combination of the SDC computing system located at the SSCL, some number of geographically separated regional computing centers and local centers at users' locations. The analysis capabilities provided to remote users should be essentially the same as those provided to users at the SSCL.

One of the most critical system requirements is that events be reconstructed at a rate that keeps pace with the recording of experimental data. The 100 Hz trigger rate and a 1 Mbyte maximum event size require a system sized for raw data sets of up to 10^{15} bytes per SSC year. Providing storage for this amount of data and access to it drive the planning of the computing system much more than does satisfying the CPU power requirements.

A major computing facility for SDC production reconstruction will be located at the SSCL (see Fig. 10-1). It will use processors that probably will be chosen from commercial RISC offerings. The output of the production reconstruction will be divided into physics oriented data sets from which physicists will draw samples of data for analysis.

Following production reconstruction, the data, or a description of the data, will be put into a database format to facilitate searches for events matching physics criteria. The ability to construct a database with the required capabilities can be inferred from present operational and developmental database systems. Continued collaboration with the computing industry and academic community will be needed to be sure that SDC needs are met in this regard.

The offline system will provide data management services for the raw event data, reconstructed and derived physics data, and data sets generated by Monte Carlo simulations. Access to data will be via a standardized query interface.

Separate computing systems that emphasize interactive capabilities will be needed for later stages of data analysis. Extensive study must be made of the methods of distributing processed data and providing access for physicists doing analysis. It is likely that both networks and physical transport of storage media will be used to supply data to a number of SDC regional computing centers (see Fig. 10-2). The regional centers will be located within the U.S., Japan, and Europe at geographically convenient places and will be repositories of intermediate size data samples.

We assume that the individual workstation, within a distributed system, will be the standard desktop device that a physicist will use both for local processing and for access to computing facilities at regional centers or at the SSC Laboratory.

Software for SDC offline computing is expected to be divided into two broad categories. The first will be a "kernel" that provides a structure for SDC software and the general services needed by all parts of the system. The kernel will be the responsibility of an SDC computing group located primarily at the SSCL. The second category is made up of the detector specific software that will be provided by the appropriate subsystem groups. The computing group will provide development systems and tools, define standards, and organize code management and distribution. They will also provide a review process to verify completed software before it becomes part of the SDC production system. Software engineering methods are expected to be used throughout, more rigorously applied in the kernel development than in the detector specific

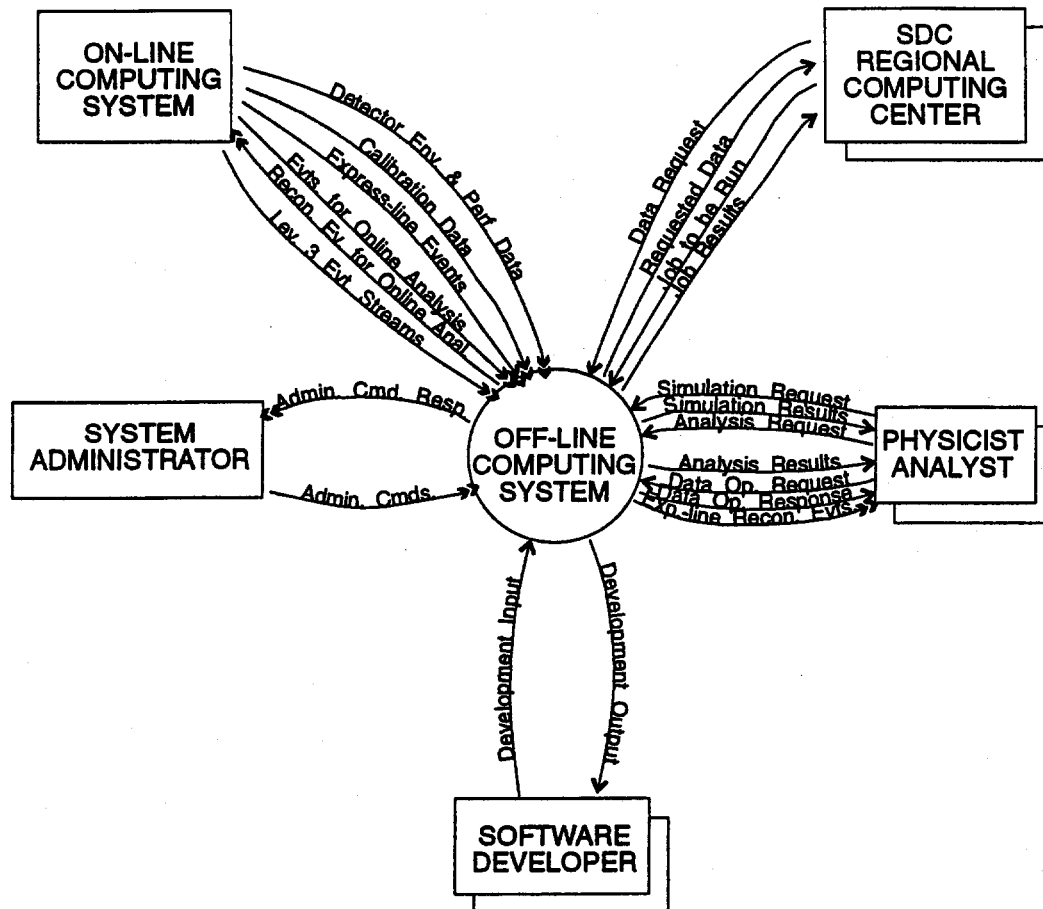


FIG. 10-1. Offline computing system interfaces.

software. Software version control and library management will be the responsibility of the system, not its users. The system will assure access by a user to the appropriate current software libraries, and will take responsibility for organizing the distribution of these libraries and tools.

The system implementation will make effective use of then-available technologies. Significant research and modeling must be performed to identify the most appropriate technologies, determine how they can best be applied, and to pinpoint possible shortfalls or bottlenecks in the system. Research is most needed in the areas of data management (database utilization, data location transparency), user/system interaction (data access, networking), and software development (methodologies, CASE tools). The focus will be on technology selection and application, not on technology development. Table 10-1 lists brief definitions of the acronyms used in this chapter.

The SDC will rely on commercial software to the extent possible. Standards defined by the computing industry will be relied on heavily to allow replacement and upgrading of software with new versions, either commercially produced or written by the SDC. The extended time scale, both before and after detector turnon, dictates the maximum possible use of standards to minimize obsolescence.

The strategy outlined above requires development of most of the kernel system before the subsystem groups begin writing the detector specific software. Thus, work on the kernel must begin essentially at the same time as the construction of the hardware components of the SDC detector. Development of the first release of the kernel system requires roughly two years, after which work can begin on subsystem software. Initial versions of all SDC production code should be complete by the end of FY1998, to allow systematic

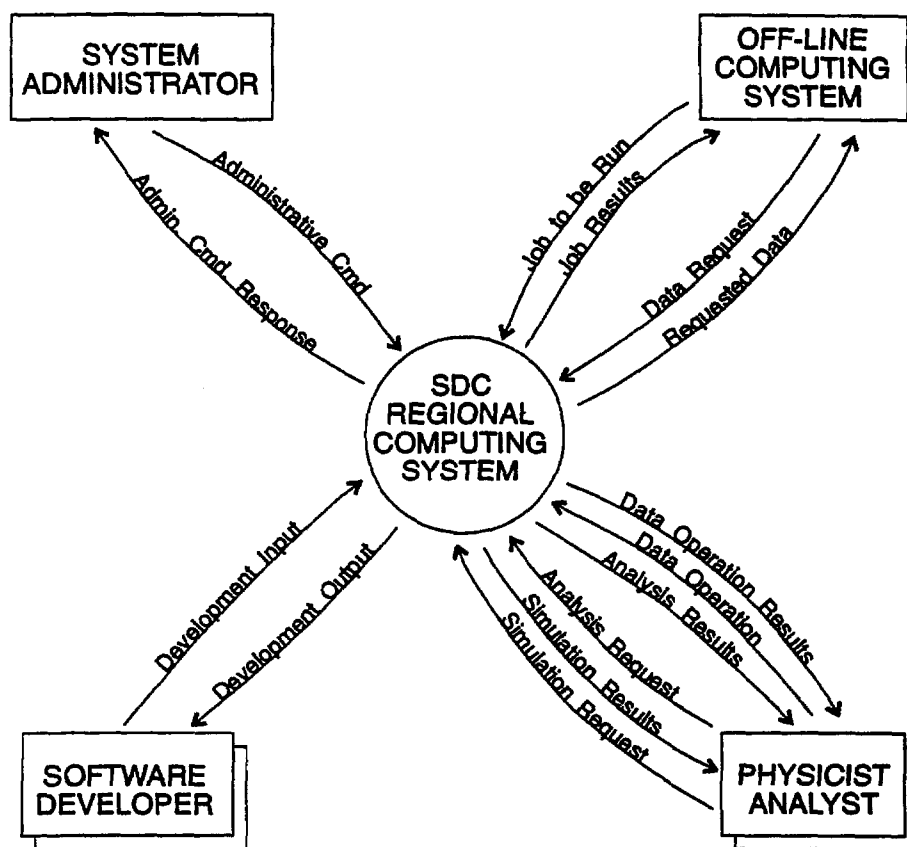


FIG. 10-2. Regional center computing system interfaces.

review and testing to ensure a capable production software system from the beginning of data taking.

10.2. Functional requirements

The first step in the creation of the SDC offline computing system is to define the functional requirements for data reconstruction, analysis, and simulation. These requirements provide the basis for software and hardware systems design. Basic SDC design parameters are summarized in Table 10-2. These parameters determine the requirements discussed below: data production, data storage management, data analysis, simulation, and communications. In addition to these quantitative areas, we have also studied our software requirements and they are discussed in the following section.

Table 10-1
Glossary of acronyms used in this chapter.

Acronym	Explanation
2167A	The formal software engineering methodology required by much of the US Department of Defense.
ANSI	American National Standards Institute. One of the US standards making bodies. Particularly associated with computer language standards.
AVS	Advanced Visualization Software. A commercial package for doing sophisticated image and data analysis while writing a minimum of code.
CASE	Computer Aided Software Engineering.
CPS	Cooperative Process Software. A Fermilab product written for the <i>batch</i> -like submission of jobs to a UNIX farm.
CMS	Code Management System. A VAX/VMS based product for code management.
CMZ	Code Management for Zebra. A code management system from CERN.
CPU	Central Processing Unit. The main computing engine of a computer.
DCE	Distributed Computing Environment. OSF's offering to allow applications to interact transparently across a wide range of networks, hardware and software platforms.
DECnet	Digital Equipment Corporation's networking system.
DST	Data Summary Tape, the output of an event processing program. Note that DST's are not necessarily on tape nor smaller than their input.
ESnet	Energy Sciences network. The US backbone network of the US Department of Energy.
FDDI	Fiber Data Distribution Interconnect. A set of hardware and software protocols for 100 Megabit/s network connections.
FIPS	Federal Information Processing Standards. Software standards which must be met for many US Government purchases of computers. Includes POSIX, among other things.
GOSIP	Government Open Systems Interconnect Protocols. The networking protocols mandated by the US Government for many of its computer acquisitions.
HIPPI	High Performance Peripheral Interconnect. Hardware and software protocols for locally connecting peripherals to computers at 100 Mbyte/s. Generally not a network in the Internet sense.
IEEE	Institute of Electrical and Electronic Engineers. One of the standards making bodies in the US. Particularly associated with the POSIX operating systems standards.
ISO	International Standards Organisation. One of the international standards making bodies. Particularly associated with OSI.
MIMD	Multiple Instruction Multiple Data. A generic term to describe parallel processing machines where each processor has its own data to work on and its own independent thread of control. Each processor in a MIMD machine may be at a different point in the program at any instant.
MIPS	Million Instructions per Second, but "instructions" is ill-defined. Rule of thumb is that 1 MIPS is equivalent to 1 VAX 11/780.
NREN	National Research and Education Network. One of the backbone networks run by the US Government.
OSF	Open System Foundation. A not-for-profit organization supplying interoperating software across a wide range of software and hardware platforms and application problems.
OSF/1	OSF's operating system offering. A version of Unix.
OSI	Open Systems Interconnect. A developing, but widely endorsed set of network protocols to allow interoperation of programs on multiple hardware platforms across a global scale.
PAW	Physics Analysis Workstation. A software system from CERN which allows physicists to interactively analyze simplified representations of their data.
Phigs	Programmer's Hierarchical Interactive Graphics System. A 3-D graphics API.
POSIX	Portable Operating System. The API's for the operating system interface being standardized by the IEEE 1003.n series of committees. It is Unix-like, but not necessarily Unix.
PVM	Parallel Virtual Machine. A software system for using a collection of computers on one program.
RISC	Reduced Instruction Set Computer, a generic term for computer architectures which are optimized for speed by using a careful simplification of the hardware instruction set.
SCCS	Source Code Control System. A code management system for UNIX, similar to the VAX/VMS product CMS.
SIMD	Single Instruction Multiple Data. A generic term to describe parallel processing machines where each processor has its own data to work on, but all processors share a single thread of control. All processors in a SIMD machine are at the same point in the program at any instant.
SONET	Serial Optical Net. Hardware and software protocols for network connects at up to 1 Gbit/s.
VAX	Virtual Address eXtension, Digital Equipment Corporation's CISC computer architecture.
VMS	Virtual Memory System. Digital Equipment Corporation's operating system for their VAX architecture computers.
WAN	Wide Area Network. A computer network on a continental scale.
WORM	Write Once Read Many. Optical disks which can be written only once and never erased. Analogous to Programmable Read-Only Memories.
X/Open	One of the international consortia recommending computing standards. Particularly associated with POSIX-level standards.

Table 10-2
Design parameters for SDC offline computing.

	Requirement
Data recording rate	100 Hz
Raw data event size (maximum)	1 Mbyte
Live time per year	10^7 sec
Raw plus processed event size	2 Mbyte
Expected DST event size	10^5 bytes
Total number of events per year	10^9
Total raw data size per year	10^{15} bytes
Total processed data size per year	10^{15} bytes

10.2.1. Data production requirements

We require sufficient computing capacity to keep pace with a 100 Hz data recording rate. The CDF collaboration estimates that an average of 250 SSCUPS-sec are required to process an event (SSCUP is a measure developed by the SSCL roughly equivalent to the performance of a VAX 11/780). This processing is dominated by the tracking pattern recognition and fitting. The SDC will need more CPU power per event. We estimate 1000 SSCUPS-sec per event for production reconstruction, i.e., 10^5 SSCUPS for the first pass reconstruction. We assume that reconstruction doubles the size of the raw data so that we will require a storage capacity of 2×10^{15} bytes for each year of operation. We estimate that we will need about 2×10^4 19 mm helical scan tapes (D-1) with a storage capacity of approximately 100 Gbyte/reel. Storage volumes are expected to be housed in an automated system.

These estimates are necessarily crude and apply to running conditions sometime after startup of the detector. Second pass raw data reconstruction will be needed at least in the first few years of operation but the event rates may not be 100 Hz during this period. In any case, we use 10^5 SSCUPS and 2×10^{15} bytes for our planning.

10.2.2. Data storage management

The expected data complexity and access requirements suggest strongly that the data be organized in the form of a database. A dedicated team at the SSCL will create and maintain the databases. The use of SDC data and its distribution requires additional database support at sites other than the SSCL. Several forms of data must be managed:

- Raw data— 10^{15} bytes/year.
- Processed data—An additional 10^{15} bytes/year.
- Selected total data samples—This is a subset of the data containing a limited number of complete events should such information be required for detailed analysis.
- Selected partial data samples—These would be datasets reflecting a users "view of the data", i.e., a selection of those features required for analysis.

It is expected that there will be two main types of data storage: reduced data sets for physics analysis on fast random access media, and the full (raw plus processed) event data sets on a long-term access medium. All data—raw, processed, calibration and detector parameters—and the relations among parts of the data will be accessed through the database. Database keyword definitions will form a metafile that will reside at the SSCL, regional centers, and, as needed, individual institutions. The distribution and organization of the database structure should be transparent to the general SDC user, whether remote or local. The database query procedure should contain useful performance information such as the probable access time, access privilege, and other penalties for the requested access.

A user accesses data through a query language by a set of keywords— event number, particle type, other physical quantities—that extract from the database only those event data segments that are relevant for his physics analysis topic. Since only the pieces of event data that are important for a particular analysis are accessed, the limited I/O resources are optimally used.

Standard DST's will be made and kept at the SSCL for general access. It is anticipated that the DST's will contain no more than 10^{14} bytes for a year of data taking—10% of the full raw data size. Specification of the number, size, and composition of required DST's awaits more detailed study of potential physics analyses. Users will be able to generate their own DST's from the standard DST's (10^9 events). Data selection for inclusion will be carried out using the same procedure of keyword database access. It should be possible to generate samples of 10^5 events for user workstation analysis within 10^4 seconds. We estimate that sufficient capacity will be required to carry this out 200 times each week.

Storage hierarchy

Online and archival storage will be based on a hierarchy of devices and media, with location-transparent access to the files. The different levels possess varying degrees of access speed, storage capacity, and cost. Total storage will support at least 2000 Tbytes per year. The highest level (fastest) of the storage hierarchy will support a 100 Mbyte/s data transfer rate for each of raw data, reconstructed data, and data used for analysis. This will require that a minimum of 1 Tbyte be stored on random-access devices at the SSCL and each regional center for short term use. A backup copy of all raw data will be archived.

File and media management

A file and media management system (FMMS) will be needed to maintain master file directories (MFD) of all online and archival files under its control. The FMMS shall follow the evolving guidelines of the IEEE Mass Storage System Reference Model, where appropriate. We require the FMMS to be able to store logically grouped sets of data as a single entity, to have at least 1000 files concurrently opened, and to be expandable to match the growth of the SDC data. As a critical resource, the FMMS must be robust and provide continuous performance; it must degrade gracefully when a device (disk or cartridge drive, operator's console, *etc.*) fails.

In order to maximize coherence within the collaboration, access to data should be location-transparent. Therefore, all data access will be through the Database Management System. File transfer shall be accomplished using, for example, the DoD Internet File Transfer Protocol (FTP) or its OSI equivalent, with the intention to be in compliance with GOSIP. All client and server functions shall be implemented. It would be desirable to perform file format conversions at the same time. The FMMS will not transfer files to a user unless it has received assurance that there is adequate space at the destination for receiving the file.

The following capabilities are needed by operations, systems, or maintenance personnel:

- ability to access, list, or modify the contents of the MFD in a special privileged mode,
- tools to monitor performance, carry out maintenance, and alter operating parameters,
- tools for recovery of data from failed media and devices, and
- tools to recovery or restore the MFD in case of corruption or catastrophic failure.

10.2.3. Data analysis requirements

Processing requirements

Many of the physics analysis functional requirements depend on the following two issues: 1) the rates for the physics topics under study and the size of the data set one will need to extract a particular signal; 2) whether the analysis will be centered at SSCL or at a remote site. Other requirements, such as user interface and analysis environment, are independent of those considerations.

Computing facilities used for analysis will fall into a hierarchy of at least two layers. At the lowest level, each physicist doing analysis will use a workstation or its functional equivalent within his local institution. The workstation should be adequate to make trial runs on events that have been selected from large samples and for repeated histogramming passes over refined data samples. Access to larger computing and storage facilities will be needed for the occasional selection pass from an intermediate-sized data set. The larger facility could be located solely with other computing facilities at the SSC Laboratory, but it will be advantageous to distribute some of these shared facilities in regional centers.

We first analyze the requirements for the user workstation. Each user will require at least a 100 SSCUPS workstation for individual analysis. On such a system it will be possible to analyze 10^4 events in 10^3 seconds at the assumed rate of 10 SSCUPS-sec for each event. This also requires that the analysis workstation system be able to read a disk at 10 Mbyte/s. Table 10-3 summarizes the physics analysis workstation requirements.

It is anticipated that a practical requirement will be for each user workstation to have 10 Gbytes of disk and 100 Gbytes of tape storage capacity. These numbers are derived from considering some potential physics analysis topics and reasonable technology forecasts for future disks and tape drives. The 10 Gbytes of disk storage will give a user the ability to execute fast passes through a sample of his data in order to tune up an analysis procedure. This would then be followed by a slower—I/O limited—analysis of the full dataset from the 100 Gbytes tape system.

The number of events to be stored in 100 Gbytes is physics-topic dependent. For some topics it will be possible to store the complete data sample with full event records, while for others it will only be possible to store samples of the data. As examples we cite the following:

1. Higgs search: For a $400 \text{ GeV}/c^2$ Higgs decaying into four leptons, we expect a full data set of about 10^5 events per year. The complete set can exist on an individual 100 Gbytes disk store.
2. Top studies: We expect a data set of about 3×10^6 events - 6 Tbytes - per year. In 100 Gbytes we can store a few percent of the full sample or about 10-20% of the summarized data.

For analysis of large data samples beyond the capacity of individual workstations, facilities will be required at the SSCL and the regional centers. To analyze one sample of 10^8 events in one day requires a system of 10^4 SSCUPS. To support several such analyses will require a capacity of up to 10^5 SSCUPS distributed between the SSCL and the regional centers.

Table 10-3
Physicist analysis workstation requirements.

Analysis Requirements	
CPU power	100 SSCUPS
Local random access storage	10 Gbytes
Local "tape" storage	100 Gbytes

Analysis environment

The present enthusiasm in the high energy physics community for PAW underlines the power of programs that provide interactive data analysis with a minimum of programming. This enthusiasm stems from the significant increase in productivity that has been experienced with its use. An important requirement for the SDC user environment is that the user be able not only to read in data and perform calculations with offline reconstruction modules, but also be able to do interactive analysis and histogramming, have the ability to incorporate personal analysis code, and to make full 3-D graphics and event displays. Such an environment will provide modules for reading raw data and applying offline reconstruction codes dynamically. We plan to study and make use of ideas from a number of advanced analysis projects being developed in the high energy physics community, for instance in the REASON project [1].

The following scenario lists the steps a user might take to gain access to the offline computing system and analyze data. It illustrates some of the highlights of the envisioned system. A more complete description of the system from the user's point of view is available in Ref. 2.

1. Environment. The user invokes the standard SDC environment on a workstation. Access to the system is gained without regard to physical location. The interface style, workspace manager, user's personal data store, file system access, and database access are all part of the standard environment.
2. Data management. Access to lists of the currently available standard data sets is provided through an SDC data browser. This makes available to the user descriptions of reconstructed event data and of "derived data sets"—data sets which were created and registered with the system during earlier analysis activities. The user selects from among these.
3. Analysis framework. Working within an analysis framework program, the user graphically chooses standard analysis modules to be applied as filters to further reduce or characterize the selected data set. The filters employ database queries that select data for retrieval by specifying conditions on its data attributes. The filters may also extract selected parameters from the full data set.
4. Programming. Analysis can rarely be completed without writing code. A set of templates is provided to embed new code modules that embody individual analysis approaches. User code is compiled and dynamically linked into the standard analysis framework.
5. Histogramming. Using a data browser, histograms, tables, scatter plots, etc., are assigned to various quantities of interest. These can come from the standard data banks or from user generated data.
6. Testing and debugging. A trial run can be made to test the event selection criteria. If necessary, the system can determine where the analysis job will be run and handle the movement of input and output data as required. If data are not local, the analysis request may be routed to the regional center where the data resides. The job is run by database engines that return to a user-specified location a test data sample or the analysis results (histograms, ntuples, etc.).
7. Cost and time estimates. After a few iterations to adjust the selection criteria, a pass through all the requested data sets is requested. The system can estimate the cost of the data access and the time to completion. The user may request that results in the form of a data set be deposited on local disk for additional work.
8. Making data public. The output data set from the above procedure may be of interest to other members of the collaboration. The system will manage the inclusion of new data sets into public lists. As data sets are less actively used, they are automatically migrated to backing storage to make room for new analysis datasets.

10.2.4. Simulation requirements

Various levels of simulation are needed to provide data to evaluate the detector designs during the period of final design and construction. A full scale simulation will verify the performance of whatever is finally built. Another role of simulation is to provide input data to test and verify the event reconstruction, data analysis, and Level 3 trigger programs. This requires simulation of real data as generated by the detector. Simulated data is processed through the event reconstruction and analysis programs like real data. In order to optimize the system and the algorithms, we need to simulate a large number of events and to process them through the entire data production chain well before first data taking. In order to improve the speed of the simulation programs, we need to investigate various techniques, including hybrid simulation programs which perform full simulation of most of the detector response but include parametrizations where possible.

We estimate the processing time needed to simulate a top event in the SDC detector to be five times larger than the approximately 10^3 SSCUPS-sec that CDF uses. Since a $150 \text{ GeV}/c^2$ top is expected to be produced at 10 Hz at standard SSC luminosity, approximately 10^4 SSCUPS are needed to simulate at the same rate. Provision for the study of many other processes of interest surely will increase that requirement by a factor of ten. Thus we arrive at 10^5 SSCUPS as the minimum required for simulations.

10.2.5. Communications requirements

SDC LAN support

The SDC experiment will make extensive use of local area networks (LAN's) at the SSCL to support the detector construction, commissioning, operation and data analysis. LAN's will be used in the detector hall and in the operations center for interconnection of workstations and various kinds of servers (data, file, terminal, print, *etc.*) and will be used in typical application areas such as detector control/monitoring, data processing, physics analysis, *etc.*

The SDC LAN will use a hierarchical design matched to the number of users, type of users, expected data transfer rates and access rates. Given the size of the SDC detector and the reliance on computer interfacing, especially during the commissioning period, the LAN will need to support up to 100 simultaneous workstations, a modest (10^4 SSCUPS) reconstruction ranch and a data archival facility. A high capacity gateway between the SDC and SSCL LAN's will be required to allow connection to the SSCL central computing facilities and high capacity gateways to the DOE ESNet WAN and other national WANs.

WAN support

Extensive use of wide area networks (WAN's) will be made for communication among collaborators and for data transfer. Because of the international nature of the collaboration, a high quality international WAN is critical. However, dependencies on the network connection into SSCL will be minimized by the existence of regional computing centers. The centers need to be interconnected by a redundant high speed WAN at 1 Gbit/s. The SDC and SSCL will need to interact strongly with the national/international network providers to insure that aggregate and peak load demands are met.

The WAN's must support current standard protocols and migration to future industry standard protocols. The WAN's must be physics-user transparent and support extended window workstation environments. The WAN's must support an optimized interactive access environment and file transfers for, typically, 200 simultaneous users from a community of 1000.

File transfer

Efficient and timely transport of SDC raw data and processed data is very important. This includes transport within the SSCL LAN and the WAN, to move data from SSCL to the regional centers and to individual physicist workstation and file servers. The transport of the largest data samples on high density magnetic or optical media will remain important. Sizes will range from 100 Gbytes to 10 Tbytes. The transfer of files of sizes 1 Gbyte to 100 Gbytes over LAN's and WAN's will be routinely generated by user requests and automated requests. The workstation environment and software systems used will require transparent access to network extended files systems of files of sizes 10 Mbytes to 1 Gbyte. For example, user DST generation of 200 samples of 10^5 events each per week requires an average network capacity of a few times 10 Mbyte/s to deliver the data to the user.

Communications aids

Communications aids for SDC can be broadly divided in to three major categories: people to people, people to computer, computer to computer.

The category of people to people communications includes electronic mail, voice (telephone and voice mail), video and text/documentation. Small group video conferencing will be used extensively. All collaborating institutions should have access to this technology at their home institution or the nearest center. The SSCL must be capable of supporting multiple independent small group video conferencing. The SDC and the SSCL should support videotaping of large meetings and conferences and distribution of video tape copies. The developing technologies of packet switch video, high speed WANs and multimedia workstations hold the promise of "personal video" conferencing. Such technology should be watched, and, when mature, added to the SDC communications support.

The category of people to computer communications includes workstation services, electronic news services, and database services. Full text and image documentation databases will be used by SDC. Hypertext access will be used, as will multimedia HELP systems and tutorials. Intelligent electronic news systems will be used extensively. Word processing matched to SDC documentation and programming standards will be important. People to people and people to computer communications must be highly integrated into the standard SDC workstation environment.

The category of computer to computer communications includes remote file access, distributed processing, database services, *etc.*

10.3. Software requirements

10.3.1. Software environments

The SDC software system must support a variety of computing environments matched to the computing tasks required for the experiment. The various software environments need to be optimized for the particular task at hand, but at the same time should share as many common features as possible to make it easy for the experimenters to move from one environment to another. Moreover, there need to be smooth transitions among the different environments, both for code that may need to be used in more than one environment and for data that flow from one environment to another. We give here brief descriptions of these environments. Later sections, in which we discuss modeling and common software tools, address the way in which the overall software system provides easy transitions among software environments. The SDC software system must support the following computing environments:

- Simulation. The software must provide an accurate model of the actual detector and physics processes under study, and the simulated events generated must be capable of passing through all subsequent stages of analysis.
- Online computing. This includes collecting and recording raw data, both event and non-event data, as well as sufficient real-time monitoring to insure the detector is operating properly.

- Production reconstruction and Level 3 triggering. These are the software programs that run on all events, “reconstructing” the raw data into physically meaningful quantities that can be used in later stages of the analysis. They require considerable amounts of computing power, and typically will not be repeated, so they must live up to the highest standards of software engineering to insure program correctness. In particular, the Level 3 trigger programs will cause events to be discarded.
- Calibration and monitoring. These jobs involve reading large amounts of raw data for specialized analysis of particular subsystems to determine calibration constants and monitor detector performance. This environment must provide easy and straightforward access to user-chosen data samples and easy ways to include personal code into the application.
- Data analysis. This environment involves the selection of events, and pieces of each event, cataloging the data samples for later use, and running physics analysis jobs on these data samples. Since data access is critical, we expect to place great reliance on database techniques, as discussed below.
- Software development. All of the codes used in the above environments need to be designed, developed, debugged and tested in convenient ways, and then moved into production environments. Coding rules and module templates need to be provided along with various development tools and methodologies appropriate for different levels of SDC code.

10.3.2. Software architectural requirements

Framework

The computing group will provide a series of shell or framework programs that will perform most of the common functions for each of the environments above and have hooks for attaching user programs. Ideally, a single framework will support all environments, with tailorable defaults to optimize for the task at hand. Tools will be provided to create and run an executable module automatically from the user's code. Techniques such as dynamic linking should be employed to lessen long delays in getting a program running.

The framework will support most application codes which process data. The framework must run on all of the platforms we support. Since the initial versions of the framework should be available before any other code development starts, we need programmers to design the framework and develop its code very early in the life of the experiment, before more extensive physics oriented software development can begin.

Software bus

Software will be written by a large number of geographically dispersed physicists and computer professionals. It will be in use for a long period of time, and most likely will still be in use after the original coders have left the experiment. For these reasons, it is critical that the software take advantage of whatever techniques are available to provide for reusability, documentation, and maintenance of the software by others.

Recent advances in software technology, *e.g.*, object oriented programming and relational database technology, have contributed a great deal to improving our ability to take advantage of the increasing performance of hardware. The concept of a “software bus” as a system, or set of standards, into which various software modules can be “plugged” is analogous to an electronics crate. The modularity this provides is extremely attractive for many reasons, none more important than it will give access to industry software with thousands of man-hours already invested.

The SDC codes will be highly modularized and standardized. The software bus framework provides well defined interfaces for application codes used for event reconstruction, simulation and data analysis. Those codes access data via the frameworks' well defined interfaces to an underlying database. The framework also supports a highly interactive user interface for data analysis and code development. Requirements on software modularity include:

- application programs may automatically be created from building blocks (modules),
- standard modules may be linked together with newly created ones to form complete jobs, and

- standard event formats, common blocks, and subroutine naming conventions shall be defined so that programs can be slotted together in a modular manner.

Standards

Most advances in software tend to be organized around standards. Although it is unlikely that the SDC will have the resources to purchase all of its needed software from industry, by taking advantage of the evolving standards we would minimize the expense of software development and be in a good position to take advantage of any applicable advances in computing that arise in the future.

Several software areas appear to be suited to some form of standardization. These are also the parts of the software system that provide the infrastructure of application independent services— operating system, graphical user interface, data management, tools, communication and networking, resource management, and distributed batch queueing.

Common tools

SDC software applications will be supported by a variety of common services. These fall into two categories: system services (*e.g.*, networking, databases and resource management) for which we expect to make use of commercial standards and more specialized high energy physics services (*e.g.*, histogramming and fitting) that we will need to write ourselves. These tools are described more fully below.

10.3.3. Data model

One of the most important tasks for the SDC software system is to provide tools that allow the user to define data structures, provide access to data stored on external media, and manage data structures in memory. We use the term data modeling to describe this array of data management and manipulation tools. This area includes tools used at software design time (data description languages, tools to model and document data structures) as well as at run time (tools for accessing, manipulating and storing data and managing memory). We briefly summarize the requirements for data modeling here.

At design time, the software system must provide a standard way of defining and documenting data structures, so that the software developers need not resort to reading through listings of common blocks. As much use as possible should be made of graphical techniques and object oriented paradigms to avoid the necessity that every user become expert in the use of an obscure data definition language.

At run time, a standard package for accessing data structures must be provided. It is anticipated that most, if not all, data structures will be stored in a database, allowing the use of powerful query languages to select the desired pieces of data. The software must provide an easy-to-use interface between the database and the user code, allowing the user to program in familiar languages while still enjoying the flexibility and power provided by the database tools. Extensions of graphical techniques used in data descriptions to the coding process would also be extremely useful. The data manipulation tools should also manage memory for the user where this function is not provided in the programming language.

Finally, the data modeling package must include I/O methods that are device and platform independent. Data should be able to be written in a common format readable on all SDC computing platforms. More efficient formats readable on a subset of the platforms may be used where performance is of paramount importance.

10.3.4. Software engineering

Software development for SSC research is a challenging task for the high energy physics community both in scale, complexity and its use of distributed computing. We must rely on the use of software engineering to manage these tasks. We need to consider: 1) appropriate software engineering processes; 2) engineering methodologies; 3) user interface design techniques and tools; 4) performance evaluation and performance-cost trade-offs; and 5) human resources.

It is expected that there will be a hierarchy of software engineering tools available for use in the SDC. This is to insure that the software kernel is thoroughly engineered while not stifling the creativity of individual physicists working on their private code. We anticipate at least four levels of software engineering:

- Kernel. The central pieces of the SDC software, primarily written by software professionals rather than physicists, will use the full set of software engineering tools as discussed below.
- Production reconstruction code. This code, mainly written by physicists, will be less constrained as to software development methods than the kernel. However, we do require that all code included in the production reconstruction be certified, which requires that it meet a set of specifications and requirements and pass a test suite to verify correctness.
- General SDC software. Use of a set of productivity enhancing software tools will be encouraged but not required. It is hoped that the tools can be made sufficiently user friendly and non-intrusive that they will be widely used.
- Private code. There will be no restrictions on individuals' private coding practices. However, there will be uniform data access methods expected to be used in all SDC code.

10.3.5. Tools

The computing group must also supply a wide variety of tools that will be used by all SDC application programs, including online applications, in a variety of environments. In the past many of these tools have been written and maintained within the high energy physics community (*e.g.*, CERNLIB). In the future, it is hoped that an increasing number of these tools will be standard or commercial packages that run on all SDC computing platforms. This requires that they be based on open systems and software standards to insure easy portability. We briefly summarize the required tools here.

Languages. SDC software platforms will support FORTRAN90 and C++. Code written in either of these languages must be cross linkable with code written in the other language. Object-oriented coding will be encouraged.

Framework. As discussed above, the computing group will provide a series of program frameworks tailored to the different SDC computing tasks. User code can be linked in to these frameworks to provide customized applications.

Data access. Standard routines must be provided to catalog and access data files residing on a variety of media and in a variety of locations, with user access being independent of the actual location of the data. Frequently used files should automatically be cached on higher performance media, with infrequently accessed files migrating to more remote storage.

User interface. The user interface should be based on a commercial standard graphical user interface (GUI), with facilities for both menu, iconic and command line input of directives. All tasks must be able to execute even from "dumb" terminals that do not support the GUI.

Graphics and visualization. Powerful tools for presentation graphics and data visualization should be available. However, the particular graphics interface should be hidden from the applications programmer by a level of abstraction. Then, changing a library will allow the code to be ported easily to any graphics standard. Code should never depend on a specific graphics interface.

Interactive and batch interface. All SDC programs must be able to run in both interactive and batch modes, with the choice being made at run time. Batch services should automatically take advantage of

available resources. The status of batch jobs, including selected user histograms, should be available for interactive examination while the batch job is running.

Database services. Standard database packages should be available and supported for both event data and parameter data (calibration, geometry, detector status, *etc.*). The database services should allow remote access, distributed query processing, transaction logging, backup and recovery, and support a user program language interface.

Histogramming. Tools to allow accumulation, display and manipulation of one and two dimensional histograms must be available. Interactive capabilities for dynamic binning and cutting (selecting events for display based on other distributions) and fitting routines should also be available.

Event display and plotting. Standard graphics packages to display individual events in a variety of formats (3-D, 2-D projections, lego plots, *etc.*) must be available. Standard features such as scanning, zoom and pan should be supported.

Function fitting. This tool provides the ability to find the best fit of multi-parameter functions using chi-squared or maximum likelihood techniques, giving parameter values and infrastructure, correlations and quality of the fit.

Parallel processing support. The use of batch “ranches” or “farms” of parallel processors should be made as convenient as possible. In particular, no compilation should be needed to change an application from running on a single processor to running on multiple processors. Rather, the determination of how many processors are to be used should be made at run time. Thus all the SDC frameworks must be capable of parallel execution. The SSCL shall provide good software support for users needing access to parallel processing resources. This support would take the form of making available software to allow reconstruction and simulation codes to run coherently on multiple machines, researching new ideas in parallel processing algorithms and architectures (in collaboration with vendors), and disseminating high energy physics standard software to run on a variety of machines.

10.4. Hardware design

10.4.1. Overview

In this section, we describe a proof-of-principle hardware architecture that could be used for the SDC. The pace of development of computing hardware is such that it is premature to design the final SDC computing hardware. We present here an outline of a system based on extrapolation of present technology to demonstrate conceptual directions and provide a basis for cost estimates (see Section 10.8).

10.4.2. Generic SDC computing model

The generic computing model of Fig. 10-3 represents an architecture that maps to emerging industry standards. This model features multiple compute servers, data storage servers, and system management servers coupled together by a high speed data network and a low speed control network. This type of architecture offers the following advantages:

- It separates communications for data and control. The higher cost systems connect those system components which need the higher bandwidth data paths. Separate, lower cost control paths prevent control bottlenecks by bulk data flow.
- It provides concurrent high speed data storage and retrieval. For example, data from the Reconstruction & Classification compute server can be stored on the Hierarchical Storage server at the same time the Analysis and Simulation compute server is retrieving data. This data movement could be accomplished via an intelligent HIPPI switch or Fiber Channel at approximately 100 Mbyte/s.
- It uses functionally coherent servers that will exploit present and future standard components, *e.g.*, ISO systems management or IEEE Mass Storage Reference Model.

- It has scalable servers, communications and storage subsystems. Each server is tailored in size and cost to match the user requirements. For example, the reconstruction and classification compute server may be several independent clusters of tightly coupled workstations attached to the network and can be easily expanded via additional clusters to meet greater future processing needs.
- It contains automated hierarchical data management, with migration and caching. For example, the Hierarchical Storage server may consist of multiple disk arrays (low latency), an optical jukebox (intermediate latency), and multiple automated helical scan tape libraries (high latency) as independent subsystems all attached to the network, yet managed via the network file server as a single integrated system.
- It maximizes the use of commercial, off-the-shelf hardware and software thereby reducing life cycle cost.

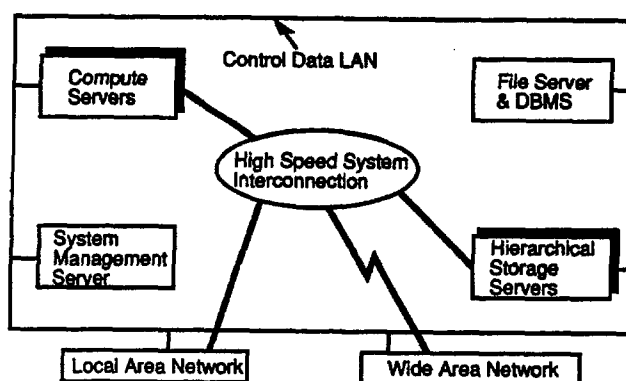


FIG. 10-3. Generic future computing model used as the template for SDC off line computing

10.4.3. Production reconstruction system

In order to guarantee the consistency of results from the reconstruction program, we anticipate that all production reconstruction from raw data will be done at SSCL. The SDC conceptual architecture diagram is shown in Fig. 10-4. Options for each major box in the diagram are summarized in Table 10-4. To construct different options for hardware architectures, we assume:

- Events are independent of each other with respect to processing. This trivial parallelism allows for cost-effective processing ranches.
- Production should keep up with the online system with a reasonably small latency. Since the production data may be used to monitor upstream data collection, we require a maximum one hour latency for some fraction of raw data to be reconstructed.
- Analysis computing resources for local SSCL users will be available. A typical analyst would require 2 to 3 jobs turned around during prime shift and several overnight. The analysis may require data retrieval from a database management system onto a local workstation disk.

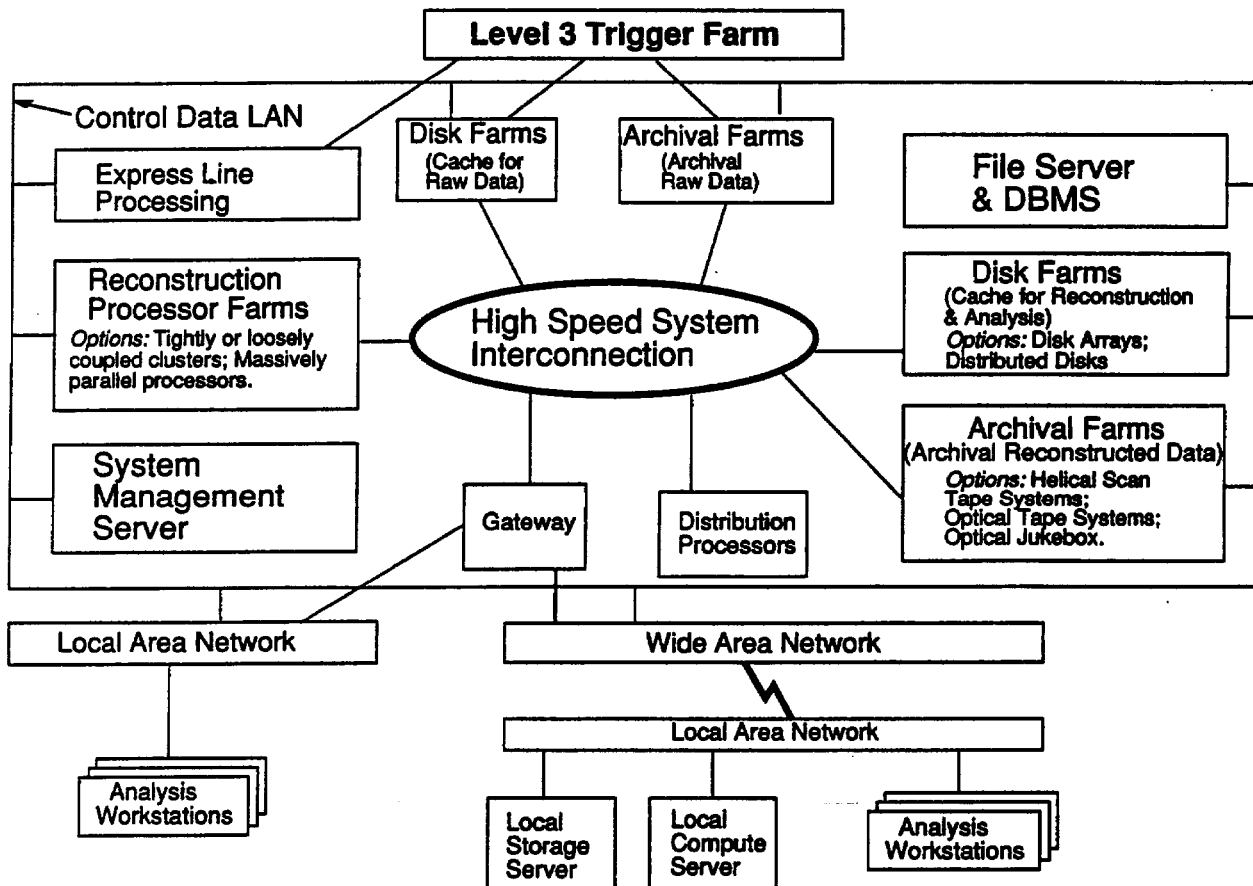


FIG. 10-4. Conceptual architecture diagram for SDC computing.

Table 10-4
Architecture options.

Category	Options
Disk Farm	High speed disk arrays Distributed disks
Archive Libraries	Helical Scan tape (Metrum, E-System, Sony, Exabyte) Future Optical Tape Libraries (LaserTape, Creo) Optical Disk Jukebox (Kodak, LSMI, HP)
Reconstruction Processor Farm	Tightly or loosely coupled clusters Massively parallel processors
High speed system interconnections	Fiber Channel standard HIPPI FDDI

We have considered three options for the production reconstruction system. All of the options below describe explicit or implicit supercomputing techniques which are multiple instructions/multiple data machines with either shared or distributed memory. Single instruction/multiple data machines, such as the CM2 from Thinking Machines, are not appropriate for event reconstruction because such reconstruction is not a synchronous lock step type of application. Also, we may wish to redeploy the reconstruction clusters to do analysis processing during times when no reconstruction is being done.

Option 1—parallel distributed computing system

This option makes use of a highly distributed and parallel computing system. This high degree of parallelism results in lower cost in hardware, a high degree of scalability, and heterogeneous device support. However, with a high degree of parallelism comes a high degree of system software complexity. This option is shown in Fig. 10-5. The processing farms consist of distributed clusters and workstations. Data access is via distributed disks attached to each cluster and workstation rather than large high speed disk arrays. Helical scan tape systems are used for archival (long latency) storage.

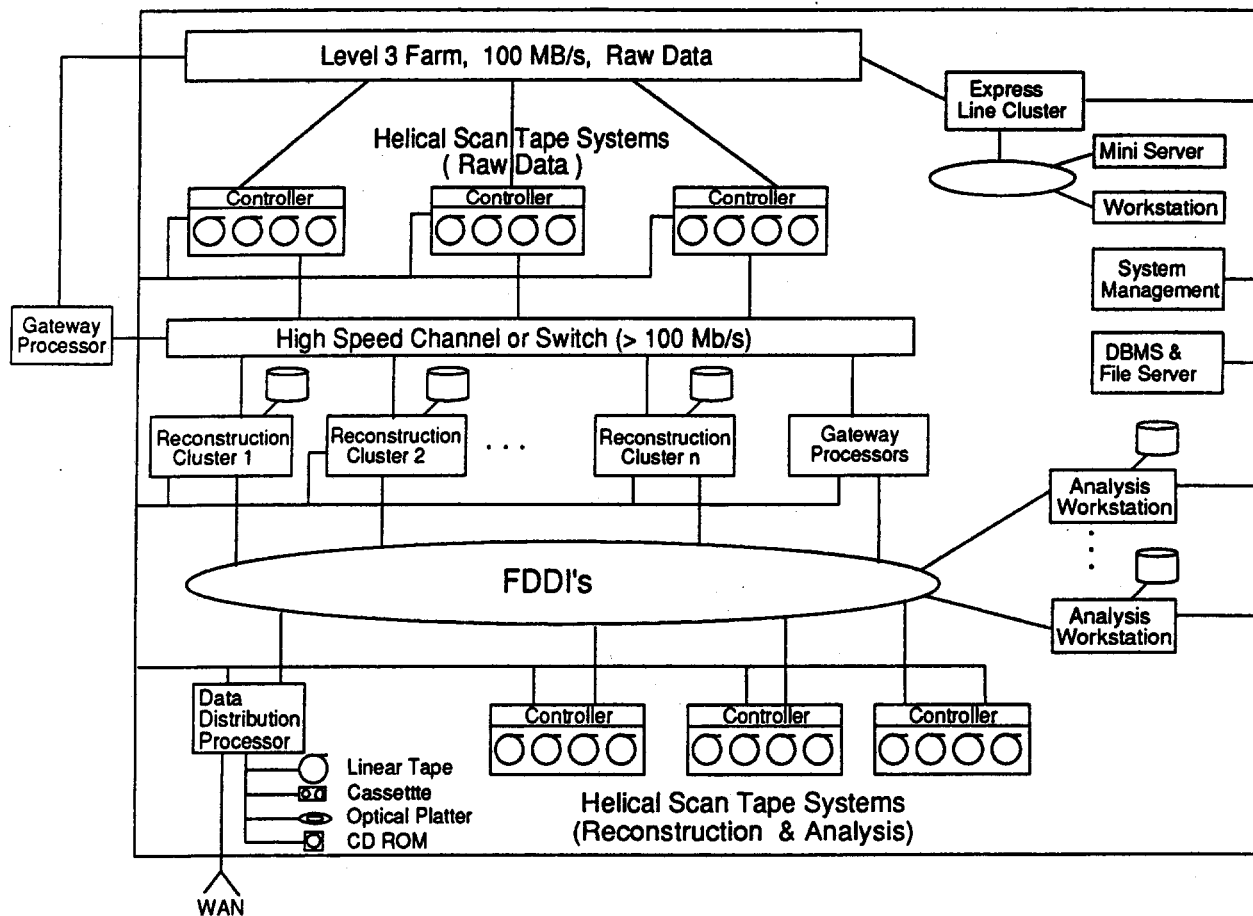


FIG. 10-5. Production computing system: Option 1

This architecture will require a complex parallel distributed operating system. As the computing industry's distributed operating system environment matures, this could be a very attractive option. Today, some prototypes already exist but they may not have all the functionality and performance that we need. Examples of those systems are OSF/DCE, Mach, Locus, NQS, and PVM. By the end of 1990s, most of the distributed system software should be commercially available.

Load balancing is easy on this architecture since the distributed environment is already in place. Local workstations are connected via FDDI, Fiber Channel, or the equivalent 1998 technology. Communication between clusters and workstations is via the same LAN as mentioned above, and there is an optional connection between the Level 3 farm and reconstruction clusters. A multi-gigabit network can be used to simultaneously distribute the raw data to both the archival systems and the reconstruction processors.

Raw data at a rate of 100 Mbyte/s are routed in the online computing system via a switching and buffering system to a series of helical scan tape systems. The helical scan tape drives are high performance,

high capacity emerging standard digital data (D-1 and D-2) tape drives which originated in the video industry. The number of drives required depends upon capacity and performance—for example, a D-1 drive can stream 40 Mbyte/s with a tape capacity of up to 100 Gbytes and a D-2 drive streams at 15 Mbyte/s with a tape capacity of 25 to 165 Gbytes. Helical scan technology has the advantage of fast data transfer and low cost—media and recording equipment cost is estimated at 0.5 cents per megabyte versus 12 to 14 cents per megabyte for optical storage. However, tapes have the disadvantage of requiring sequential access for retrieval versus direct access to data with an optical platter. Furthermore, since helical scan drives do not perform well if tape motion is sporadic, the online computer system must spool Level 3 data to disk and then to tape to keep data flowing in a continuous stream. Also, depending on the vendor, the systems management software would participate in managing which drives are used for reading, rewinding, and writing to the reconstruction clusters. Tape storage in the archive could be manual or automated via robotics.

Production reconstruction is handled by a group of small clusters of processors. Each cluster contains a small number of coupled CPU's, *e.g.*, 32 CPU's connected via a high speed backplane. A cluster manager schedules and controls routing of data from the online computing system to each CPU. There are no data communications between CPU's in the cluster; a single CPU processes a single event. If we assume that 1000 SSCUPS-sec are needed to process a single event, it will take 10 seconds to complete on a 100 SSCUPS CPU.

A hierarchical, industry standard compliant mass storage system will be used. Such a system will provide for the long latency storage, medium latency storage, caching, and file migration of event data. A database manager will be co-resident with the name server of the mass store system, which has knowledge of file names and directories. The database manager has local storage of "metadata" which facilitates selected retrieval of data from the data storage system. The name server and database manager server are connected to the low speed communication LAN since their job is primarily to direct the transfer of data between analysis workstations, reconstruction clusters, and storage subsystems. A systems management server which would comply with the ISO standard will maintain the integrity and consistency of data files and databases.

Option 2—parallel distributed computing with centralized disk arrays

This differs from option 1 in that centralized disk arrays are used to buffer data for reconstruction, analysis and simulation. The advantages of this option relative to option 1 lie in the reliability provided by high speed switching of a failed tape drive or cluster. The aggregate throughput of the disk arrays needs to be greater than 400 Mbyte/s (200 for reconstruction, 200 for analysis and simulation requests). Helical scan tape systems are used for the archive libraries; small clusters are used for reconstruction processor farms; and an optional disk farm is provided for direct caching of raw data from the online computing system so that raw data tapes do not need to be hand carried from the on-line site to offline site. A fiber optic link can connect the Level 3 farms to the offline disks. The fiber needs to have a bandwidth of at least 100 Mbyte/s in order to keep up with the Level 3 data rate. In this case, the raw data are sent to both archival systems and disk arrays simultaneously.

Option 3—supercomputing

Option 3 is identical to option 2 except in the use of a single cluster of processors for reconstruction. This single cluster would have to provide 10^5 SSCUPS to keep up with the Level 3 event rate. An example of this single processor cluster is the Sigma Touchstone supercomputer from Intel, rated at 10^5 SSCUPS, and the Vulcan Research Machine at IBM Yorktown. We note that at this date, a number of private supercomputer companies are working on multiply connected massively parallel computers with shared memory based on superscalar RISC processors and high speed interconnects.

Option 3 suffers from the fact that these "supercomputers" have built into them circuitry which minimizes the complexity of parallel processing (shared memory and processor-to-processor communication

for example). Such circuitry is not critical for the SDC reconstruction, and results in higher costs. However, the system management of option 3 would be simpler than that of the others.

10.4.4. Simulation system

It is possible that simulation data production can be distributed among collaboration institutions. In this architecture, we need to install CPU power sufficient to handle both simulation production and reconstruction at each simulation site, and store the simulated data in a manner that all can access. Also we need to guarantee the consistency of results produced at all locations. This requires version management of programs, simulation parameters, *etc.* These facts imply that some aspects of the production of simulation data would be simplified if done at one site—the SSCL. Even were mass production of simulated data done at SSCL under definite coordination, each institution would want the capability to do simulation of processes of special local interest.

10.4.5. Physics analysis system

A reduced sample of data will be maintained at the SSCL for fast access and frequent filtering of data in order to facilitate repeated physics analysis. A data size of 10^{14} bytes (10% of the raw data size) would contain all of the events from a year's running with a compressed or reduced representation of the parts of the data needed for most physics analyses.

Facilities should be provided at the SSCL for a portion of the physics analysis to be done by the SDC. Computing at the physics analysis stage involves both batch and real time processing (visual event scanning, interactive analysis tools, *etc.*). By the time we start physics analysis of SDC data, wide area networks of 1 Gbit/s and international networks of 100 Mbit/s are expected to be available. Such high bandwidth networks may enable us to centralize all CPU resources needed for analysis at one site. However, we must be careful of this approach for the following reasons: 1) the effective bandwidth of a network is always much smaller than its maximum; 2) even if a network has infinite bandwidth, there is always a round-trip time delay which causes serious deterioration of performance of real time processing especially at a site far from the SSCL, and 3) we would need to maintain high bandwidth networks between SSCL and all institutions in the collaboration. The highly international character of the SDC makes it imperative to plan for computing resources located outside the U.S. Furthermore, the very large number of U.S. physicists in the SDC, geographically dispersed throughout the country, will require substantial computing resources outside the SSCL. We have included regional computer centers in the SDC architecture design. These centers would exist at a few institutions in the collaboration as shown schematically in Fig. 10-6. In this architecture we minimize dependencies upon the network connections into SSCL, and allow for the highest possible network bandwidth to exist between a center and nearby institutions. We would be utilizing the resources which already exist in the computer centers of major institutions in the collaboration.

The number of such centers is still to be fixed. However, more than one such center will be needed in the U.S., one in Japan, and possibly one in Europe. A regional center would have the processing and storage resources to support repeated analyses by many active users. Each center would have about 10^4 SSCUPS of processing power. The physicists who use these regional centers will be located at their home institutions. Each of them will have a workstation able to display with X-windows or a successor.

10.4.6. Communications

To meet the requirements for the SDC LAN the basic hardware technologies will include combinations of 10 Mbit/s Ethernet, 100 Mbit/s FDDI, and 1 Gbit/s Fiber Channel or SONET. Industry and international standard protocols will be followed, starting with TCP/IP and Decnet and migrating to the OSI suite. The SDC LAN will be designed jointly by SDC personnel and SSCL network staff, and installed, maintained and operated by SSCL network staff.

WAN capabilities may be provided by ESNet or other national/international scientific networks, in which case the regional centers and the SSCL should be backbone sites. University based and other remote institutions will obtain WAN connections and services via a range of organizations providing network service

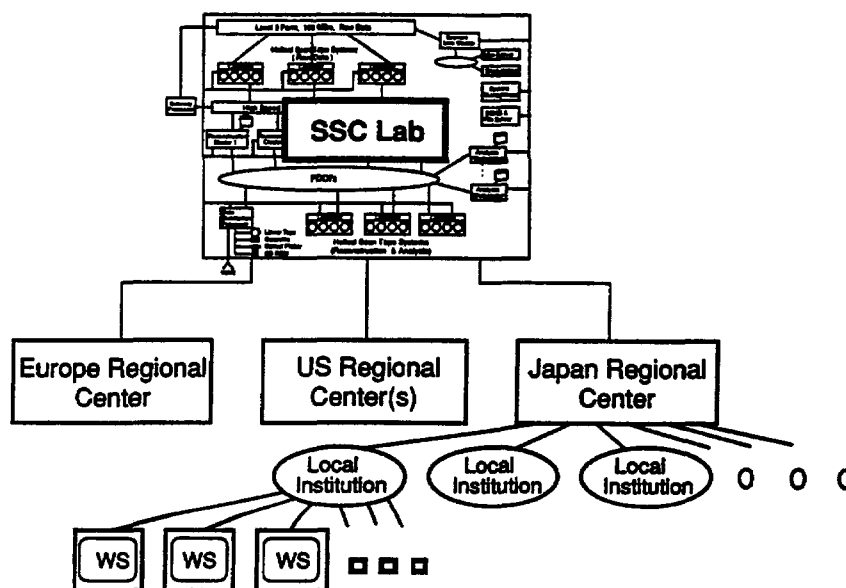


FIG. 10-6. Schematic view of SDC computing hierarchy with a system of regional computing centers for SDC analysis. Paths are shown for each regional center to get data from the central system at SSC Laboratory. The regional center in turn provides data and computing cycles for a group of local institutions (shown only for one center) at which users interact with the whole hierarchy via the workstations on individual desks.

providers, *e.g.*, ESNet, direct connection to regional centers, regional academic networks, *etc.*. The SDC and SSCL networking groups will need to provide coordination for the design, management and allocation of resources in the WAN's.

10.4.7. Hardware design verification

Any hardware design proposed must be able to satisfy the performance requirements described above. We have begun a program of computer simulation of the behavior of the proposed offline system. We are using the RESQme system, on loan from IBM, which is designed to simulate networked computer systems and allow rapid construction of sophisticated models and automatic gathering of model performance data.

To date, we have built models of the offline system in both centralized and distributed forms. The centralized model has all data and processing concentrated at SSCL, while the distributed model has five regional centers caching the data and sharing the processing. The models reflect contention for tape robot devices, processor resources (CPU and disk I/O are not distinguished in the current models) and network bandwidth. We have simulated the models with loads that represent the requirements described earlier, and at higher loads. Details of the models appear in Ref. 3.

Conclusions at this stage are very tentative. Both the centralized and regional models appear to handle the proposed load with enough reserve. The nearest bottleneck seems to be contention for the tape robot devices. Because the transfer of large data sets from the tape drives takes about an hour, short jobs can back up significantly. One possible approach to help the response of the system is to allow the long tape transfers to be interrupted by a short job and restarted later. This requires assurances that the rate of short jobs doesn't become so high that long jobs get starved for access to the tape robot.

Our immediate plans are to revisit the hardware requirements to specify them more precisely, to refine our models to more accurately reflect the behavior of the various components, and to study the data access patterns of large experiments such as CDF. All this information will be folded into the models for analysis.

10.5. Software design

Offline computing covers two different bodies of software. One, written by individual users, contains code specific to particular physics analyses. The other is more generic in character and includes event production reconstruction codes, event databases, and the overall framework within which users do their analysis. Since it has wide impact across the collaboration, the latter set must be of the highest quality obtainable and should hide differences in hardware and software platforms from the user's analysis code. Much of this section deals with problems faced by the developers of this generic code. A successful implementation will mean that most collaboration members can ignore these issues during their analysis activities.

There are two dominant requirements. For production, the system must, on average, keep up with the data taking rate. For analysis, the reconstructed data must be made available in a timely fashion to users. Both of these requirements are met by the use of loosely coupled multiprocessing. In the case of production, we use the now-standard processor farm. For analysis, we assume the basic hardware platform will be geographically extended to include regional centers and users' workstations.

We make the assumption that data analysis will be done within a database model of computing. That is, users' programs will query the database for the relevant event data fragments only for events of interest. Such data fragments are the information banks stored in the database. The reconstructed data by themselves are expected to be too large to store on random access devices, so we must also assume that the data can be modelled and structured in such a way that most of it is accessed only rarely and can be migrated to magnetic tape backup store. As an example, fitted track parameters need high availability, while track hits can probably be archived since they are seldom accessed. It is certainly not clear at this point that database technology will be sufficiently powerful to allow all reconstructed data to be in a database. However, we do expect that everyday analysis will be done on data samples in the form of a database.

To avoid continuous rewriting of code as it migrates from analysis to production to Level 3, a single framework interface definition will be used for all these applications. A variation of this framework will provide for adding Monte Carlo modules as complex as full detector simulation of raw data or as simple as directly loading analysis data banks.

SDC software will be in use for as long as twenty years. During that time we can expect dramatic changes in hardware and software technologies. To allow us to insert these new technologies as smoothly as possible, we will design our software to use the broadest and most widely accepted standards.

10.5.1. Architecture

Software architecture overview

A discussion of the SDC software architecture design goals is given in Ref. 4, Section 4.1. The resulting architecture is summarized in Fig. 10-7. This design calls for isolation of system components by explicitly defined application program interfaces (API's), which detail the syntax and semantics of library routines. There is a minimum of cross-coupling among the various modules to allow relatively painless replacement of modules as standards and technologies develop.

Most SDC collaborators will only see the highest level API which describes how their analysis routines fit into the common framework and provides libraries for common analysis activities such as histogramming and event selection. Note that each user's analysis modules may eventually appear to another user as part of the top level API, *e.g.*, tracking should appear as a common service to someone working on calorimetry.

Since the top level API is highly specific to SDC it will have to be custom developed. However, there are many common services needed to support this framework. These include graphical user interfaces (GUI's), network services, data management services and system management services. These services are widely recognized as fundamental to the whole computing industry and are moving quickly toward standardization. We will gain considerable leverage by using standardized versions of these API's.

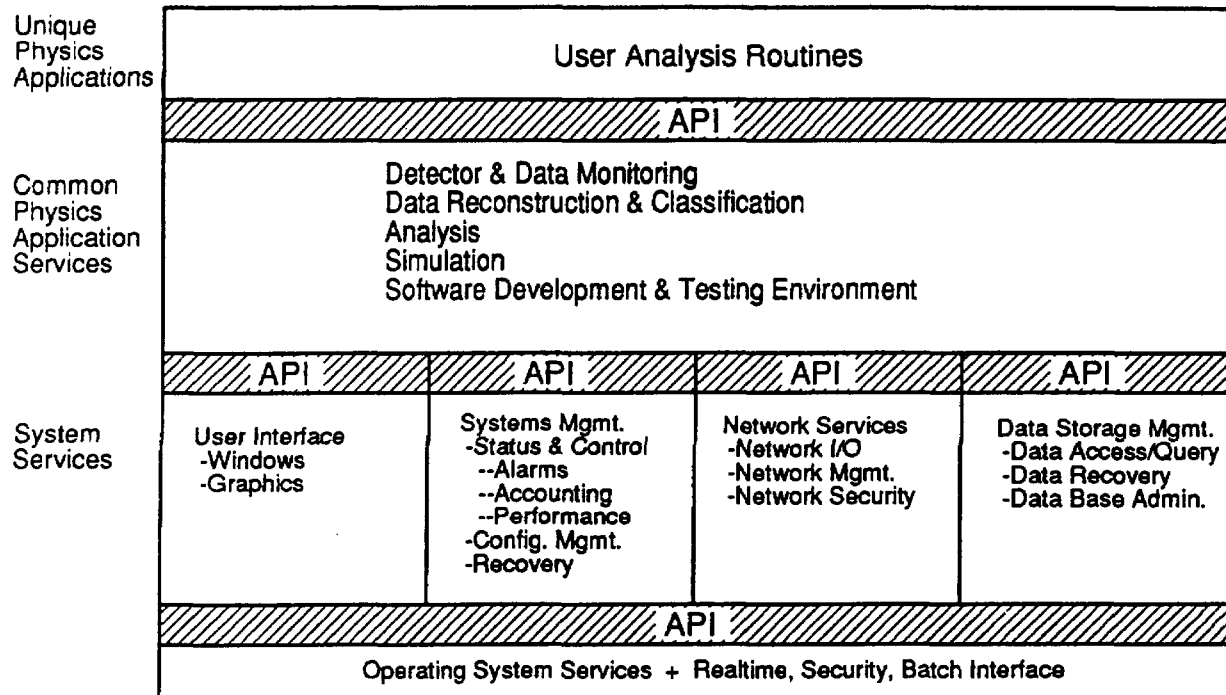


FIG. 10-7. Modular software architecture model, showing layers of software connected by well-defined application program interfaces (API)

Underlying everything is the API for the operating system. Use of a standard API like POSIX rather than a particular operating system, will allow the whole SDC software system to run on UNIX, VMS, or any other POSIX-compliant system. We can continue to utilize the installed base of VAX/VMS machines and be in a strong position to utilize whatever technological breakthrough may supplant RISC architectures.

Software infrastructure/framework

Operating systems/portability

For the fundamental framework that can run on any collaborator's workstation, the code will conform strictly to the POSIX standards. As currently defined, POSIX standards are not sufficiently complete to handle some of the more sophisticated multiprocessing servers that will reside at SSCL or the regional data centers. For the cases where portability requirements are much less stringent, we follow the OSF/1 precedence list of applicable standards. In particular, we will use POSIX/FIPS requirements and then X/Open standards.

In some cases, necessary services are not yet provided by these standards. For immediate, exploratory system development efforts, this implies that choices will have to be made before official standards are in place. The choice fundamentally is between OSF/1 and AT&T's System V Release 4. Both of these should be increasingly available in the next year and may effectively merge in a few years (IBM's new AIX Version 3.2 supports both the System V Interface Definition Issue 2 and the OSF Operating System Interface). If a choice must be made, we prefer OSF/1 as a baseline operating system because of the availability of the Mach kernel primitives for light-weight threads and user-mode virtual memory managers.

Parallel processing support

For multiprocessors, such as the production farms, no commercial product or standard is available at this time which meets our specialized needs. Therefore, we will build on the Fermilab application CPS written to use common networking services. The OSF/DCE system, now available for testing on the IBM RS/6000, may provide sufficiently potent services to supplant CPS. For multi-tasking systems, such as multi-threaded data servers, we will use the POSIX light-weight thread package.

User interface

User interface standardization plays an important role in the formulation of the production, simulation and analysis processes for the SDC off-line computing and is an important aspect of software portability. Since SDC offline computing applications will be distributed over a heterogeneous environment of computer hardware, software, and networks, the user interface facility must provide the flexibility to allow the user to interact with programs within such an environment. Standards transparent at the user interface are appropriate since they provide a foundation upon which to build a desirable user interface. A consistent user interface across all systems on which the application resides promotes ease of use and simplifies training.

We assume that the X Window system will survive and continue to evolve. We adopt it as the basic graphics framework. While the choice of window manager should properly be the user's, the current state of the art does not allow programs to be easily written to conform to multiple GUT's. It is possible that this situation will change, but for now we choose to follow the evolution of the Motif GUI. Starting with X11 Release 5, Phigs Extension to X (PEX), will be available and becomes the natural 3-D graphics API for event displays and the like. The application framework will capture common interaction activities (event displays, histogram manipulations, menu building, etc.) so that most users will not have to deal with X/Motif at too low a level.

Interactive and batch user interfaces

Users can submit batch jobs without needing detailed knowledge of what resources are available at that particular site and time. Distributed queueing software will insure that the batch job is directed to the available computing resource that will minimize turn around. Since the Network Queueing System (NQS) is widely available and forms the basis for the POSIX committee's deliberation on a batch interface, we adopt it. No method is known at present to provide effective batch interfaces for X windows based programs. We will probably have to develop our own script language, although recent work on a macro facility for X may help [5].

Systems management

Systems management is a complex problem because of the large number and wide distribution of processors involved. We seek a system whose target machine maintains its own information about the software versions on it so as to avoid a need for a centralized system. The system also should allow software to be pushed from the source or pulled from the target. The system should be scalable from single machines to department clusters to the entire SDC enterprise. OSF/DMS (Distributed Management System) is designed to meet all these requirements. It should become available within the next year. We will evaluate it at that time.

Network services

We assume that, for levels below the transport layer, it will be a technical exercise for the networking experts to connect DECnet, TCP/IP, OSI, *etcetera*-net together. The real issue will be interoperability and networking within a heterogeneous system of computer hardware, software and networks. This involves issues of data formats, resource/server name resolution, Remote Procedure Call registration, *etc.* OSF/DCE appears to be the most broadly accepted, most complete and most operating-system-independent candidate. Reference implementations are available now for the IBM RS/6000 (AIX operating system) and DEC 3100 (OSF/1 operating system). An implementation for Intel-based machines running System V Release 4 will be released within six months. We will evaluate this solution over the next year.

Data storage management

While we have decided that the reconstructed data will reside in a database, it is not yet clear whether that will be a relational or object-oriented system. This question is being studied under the aegis of the High Performance Computing and Communication Initiative. A decision should be made this year. In any case, the many other items—calibration constants, hardware inventories, replacement records, *etc.*—which should be stored in a database can be handled adequately with a commercial relational system.

We assume that the vast quantities of data we deal with will be contained in a mass storage system compliant with the IEEE Mass Storage Model.

10.5.2. Development process

Software engineering and other tools

Automated software engineering tools greatly enhance the productivity of programmers and the quality of their work. However, such systems can be quite expensive and the SDC may not be able to afford to supply the tools to all collaborating institutions. Our current plan is to provide commercial tools to the principal developers of the software infrastructure and also make the tools available for general use from the central laboratory and the regional centers. This will likely prove awkward for the more casual analysis tasks. General users will be encouraged, but not required, to obtain compatible tools.

While not required, we expect most of the software infrastructure code to be written in C++. We take the current Version 2, as defined in Ellis & Stroustrup, as the basic language and restrict ourselves to that. We will adopt the ANSI standard version when it becomes available.

We expect that most of production reconstruction code, simulation code, and general physics applications will be written in FORTRAN90. We take the ISO/ANSI standard as the basic language and restrict ourselves to that. Since all recorded events will be run through this code, it will be held to a higher standard than normal user code. As a consequence, it will be required to be incorporated into the common software code management system and the validation and verification system. The developers and integrators of this code will have access to sophisticated analysis, design and debugging tools.

Physics analysis code will probably be written in FORTRAN90. Much of this code will be written by individual users with the tools they have available at their home institutions. These will include application development tools written for SDC, but may not include some of the more expensive commercial systems available to the production coders.

Code management

Physicists writing their own analysis code may elect to use a code management system. Common application services and production code *must* be rigorously kept in a code management system. Since these codes need a centralized repository, their management may utilize a comprehensive and relatively expensive system which need not run on every workstation, but would be available in a client/server mode. Since individual analysis code is usually the basis for new production code, it would be desirable to provide a subset of code management tools that would be attractive to physicists for most programming projects.

We do not yet have enough experience with various systems to make an informed choice. There are efforts underway to place the simulation shell (Section 10.6) into *sccs/make* as well as CMZ to evaluate their relative merits. We will also work with the IBM Independent Research and Development (IRAD) project to build a coherent CASE environment to understand available products. We will also look into using an object-oriented database as the code management system. Our goal is to make a decision on a code management system in 1992.

For the distribution and installation of code to be a practical operation, we assume as a minimum the existence of *sccs* and *make* on all potential target systems.

Data and system analysis

There are several data and system modeling approaches available. It appears that object oriented approaches can add significant features, such as better module decoupling and reuse. A review of the most popular methodologies appears in Ref. 4, Section 6.4. Some preliminary object modeling work in the physics domain has been done by IBM. A summary is given in Ref. 4, Section 4.2. Since a high level of abstraction is required for a successful environment, similar models for the analysis framework will be started next.

There are various tools available to help support the analysis process. We are currently examining these with the help of IBM. We will recommend a methodology and tool set by the end of 1992. We do not expect this particular tool to be the only option for the lifetime of the experiment, but we do expect the use of formal methods to last that long.

We will require the use of the formal methodologies for the framework programs and common application services. In addition, code which may have been developed during analysis but which needs to be incorporated into the production or Level 3 systems will not be accepted without the supporting documentation, *etc.* of the design methodology.

We find Ybos, Zebra and other traditional FORTRAN based systems inadequate for implementing the data model in a fashion compatible with the assumed database model of computing. These older tools lack sufficient power to declare new user data types and to retain at runtime the structure information needed for data browsing. Also, many of their memory management and data structuring tasks have been subsumed into FORTRAN90.

Validation and verification

The groups responsible for building the production and Level 3 programs and distributing the standard libraries will need automated tools to help do regression tests of the software as algorithms mature. Such tools would automate the process of running tests against standard data sets, comparing the results and calling out discrepancies. Note that this implies that the developer must provide test datasets and expected results. Such a system exists under VMS as DEC/TMS (Test Management System). We are searching for a comparable system under UNIX.

10.5.3. Software sources

Using a highly layered software architecture with well-defined API's (see Fig. 10-7) allows the various libraries to be obtained from the most appropriate sources. The following sections list those sources and the codes to be obtained from each.

Common commercial software

For maximum reliability, off-the-shelf commercial software will be used for the lowest level application code. Much of this is operating system level code which will be obtained from the hardware vendors and includes system services, distributed computing, user interface, systems management, batch services, and as much of the database computing system as possible. The rest will come from third-party vendors. This includes software development and testing environment, including compilers and sophisticated debuggers, such as ObjectCenter (formerly Saber-C++).

Custom professional software

Code at this level is too specific to be available off-the-shelf, but should be of comparable quality since it will be widely depended on by the collaboration. It may be developed through contract with a software development company or by the professional staff at the SSCL. Software categories for which this approach might be appropriate include physics analysis framework, simulation framework, detector and data display, and other software development and testing environments for physics applications (*e.g.*, a customized version of AVS).

Common high energy physics software

There is a large body of reliable code for high energy physics most of which has been developed at CERN. As much of this as possible will be reused. However, several of these libraries are elaborately interconnected. In order to maintain design freedom, we do not commit to support of all these libraries at this time.

Physicist software

This level includes all the code which is particular to the SDC. For the vast majority of the collaborators, this is the level at which they will program. Software categories include:

- Data reconstruction and classification.
- User analysis routines.
- Problem-specific simulation.

10.5.4. Software development process

This section investigates software engineering methodology and life cycle options. A detailed discussion of this topic is found in Ref. 6. The software development process must be tailored to different classes of software in SDC offline computing: kernel, production reconstruction, general physics utilities, and private user code.

Methodology options

A methodology must cover the entire software engineering life cycle, which includes analysis, design, code, test and maintenance. Three basic methodology options are considered : structured-analysis/structured-design (SA/SD), Jackson structured development and object oriented. These represent three different analysis and design styles which result in different implementations. Contrary to popular belief, structured and object oriented methods have evolved in parallel since 1966 and have much in common. For example, SA/SD and object oriented methods support three orthogonal views of a system:

- Object view. Entity-relationship diagrams for SA/SD; class-object diagrams for object oriented which portray class attributes, functions (services or methods) and class associations including composite and inheritance relationships.
- Dynamic view. Usually some form of state transition diagram.
- Functional view. Usually data flow diagrams and structure charts.

SA/SD emphasizes function first and tends to be more top-down. Object oriented emphasizes data (objects) first and tends to be more bottom-up, although in practice it can be a combination of bottom-up, top-down and inside-out. SA/SD separates function hierarchy from data while object oriented has a single object hierarchy with data and encapsulated functions.

Structured-analysis/structured-design

Procedures and algorithms are uppermost in the designer's mind. Initial analysis is a statement of "what" the system is to do. Design is a statement of "how" it is to do it. Typically, high level functions are decomposed using data flow diagrams and then structure charts. Data flow diagram bubbles represent data transformations and the data flow between bubbles represents interfaces. Data stores are also modeled. Data flow and stores become the data dictionary. Databases are modeled separately as entity-relationship diagrams. Realtime concurrency and sequencing can be modeled with state transition diagrams. Resultant implementation packaging is function oriented (procedures).

Jackson structured development

Jackson uses some object based concepts (data structure methods) to initially model the real world. However, it is essentially functional decomposition where the data structures are used to assist in functional decomposition. It uses complex pseudo-code to model realtime sequences and concurrency.

Object oriented

The object oriented paradigm combines data and procedures, but deemphasizes the procedures, stressing instead the encapsulation of data and procedures together, exemplified by the clear and concise specification of the module interface. The system is viewed as a collection of object classes which encapsulate data and services. Classes may communicate in a client-server fashion by passing messages. Class hierarchies are modeled, including aggregate (whole-part) and superclass-subclass (inheritance).

Class modeling during analysis and design focuses only on the externally visible part of the class's attributes and services. Hence, the focus is on data and service abstraction. The details within the class are hidden until late in design. Detailed design, within a class, can utilize data flow or Buhr diagrams to model services. Some variation of state transition diagrams is used to model high level class concurrency and, later in the design, concurrency within a class.

Recommendation

The object oriented style is recommended for analysis, design and development of the kernel software including the physics analysis framework. The object oriented style would be highly beneficial if used for augmentation or reengineering of any of the common physics applications software. An object oriented style is not really necessary for user code since it will be integrated within an object oriented framework. End users can elect to use an object oriented language such as C++ or a modular oriented language such as FORTRAN90. The object oriented style focuses on software modularity around data objects. This should result in a software architecture which is easier to evolve and maintain because the software modules map more closely to real world objects which change. This is particularly important because the SDC software is expected to evolve through prototyping and simulation.

Life cycle options

Life cycle options address the software planning and procedures required for cradle to grave software construction. This is generally referred to as the software development process. It is important that the software development process be defined, documented and inspectable. Collaboration members must define and assume ownership of the process. The process can be a living document which can be modified throughout the life cycle. The basic options are waterfall and spiral.

Waterfall

The waterfall life cycle, *e.g.*, the military standard 2167A, steps through analysis, design, code and test with discrete deliverables and reviews at each major milestone. The process is usually tailored to meet the needs of a particular program. System, hardware and software development plans are produced. The analysis stage produces top level requirement specifications and operational concepts culminating in a systems requirements review. High level design produces a systems design specification and a systems design review. At that point functional subsystems are identified for both hardware and software. Subsystem requirements are documented and reviewed. Subsystem designs are documented and reviewed typically at a preliminary design review. Detailed designs and test specification are developed and reviewed. Actual coding typically does not begin until after the preliminary or critical design reviews are completed. Software is released in increments with systems tests scheduled along the way. At the end some form of formal acceptance testing is performed before final delivery of the product.

In practice, the waterfall allows for some cycle back into the previous levels, *e.g.*, to update the requirements or design based on changes made at the lower levels. The waterfall has been used in conjunction with early prototyping in area such as user interfaces and realtime performance.

Spiral or fountain

The spiral or fountain model takes the same basic steps of the waterfall model but repeats them in small cycles. It is particularly useful for evolutionary development.

Preliminary recommendation

The process used will be the subject of study in 1992. At this point the spiral type of model seems more suited for kernel, framework, and physics application software development because of its evolutionary nature. Also, kernel software with its potential high commercial content would facilitate early prototyping in combination with reused and/or augmented framework and physics application software.

10.6. Detector simulation

Simulation of both physics processes and the detector response to them is crucial both as the detector is being designed and as it is taking data. A unified simulation program is being developed for use in design of the detector and in preparation for data analysis. In this section, we describe this program, SDCSIM, and outline future plans for development of simulation software.

10.6.1. SDCSIM goals

The development of SDCSIM was launched in late 1990 with two objectives. First, it should provide a common environment for simulation work for detector design and design verification. The program should include event generation, detector simulation, event reconstruction and event analysis. The codes should be developed by each group proposing detector components and integrated into one program so that a fully integrated detector system could be studied. Second, it should provide a unified software system to which many people could contribute from a very early stage of the collaboration. From past experience, we knew it could take some time to design and install a formal software system following principles described earlier in this section. Therefore a simpler approach was adopted to produce a program that would be functional as early as possible.

10.6.2. Software design of SDCSIM

In order to satisfy these goals, we established the following design criteria for SDCSIM: the system should be simple, modular, platform independent, use existing software where possible, and make maximum use of standards. Based on these criteria we chose the following for the system:

- Language—FORTRAN77 with almost no extensions.
- Partitioning—Programs are divided into many small packages with communication between packages done only through well defined data structures.
- Data access—Simple home-made interface from each package to internal data structure. The internal structure is maintained by ZEBRA from the CERN programs.
- Code management—PATCHY with minimum functionality.
- Detector simulation framework—GEANT.
- Shell—Overall program framework is a simple, home-made framework that binds various program packages and provides an interface to the data structure.

10.6.3. Current status

Interface programs to event generators and to GEANT, and to various detector simulation and event reconstruction packages have been developed by each detector group. Both the framework and the SDC program library are supported by the computing group at the SSC Laboratory.

Three commonly used event generators, HERWIG, ISAJET and PYTHIA are available and the detector geometries of the major technology options have been installed for detector simulation in GEANT. The detector simulation packages now output digitized signals or hits in tracking detectors and in the muon system. Energy deposit in the calorimeter is also reported with fast or full shower simulation. The reconstruction program finds segments in each tracking detector and does track fitting after linking segments to form a track in the central tracker. Linking tracks between the central tracker and the muon system is under development.

The program library has been established on a VAX (SSCVX1) and distributed to many institutions within the SDC. The program will run under VAX-VMS and on various UNIX platforms. It will also run in a parallel mode using CPS and in a batch mode on the Physics Detector Simulation Facility at the SSCL.

10.6.4. Future plans for simulation programs

As outlined earlier in this chapter, the major task of the computing group during the remainder of FY1992 is the design of the core software system for SDC that will provide the structure and central services for the broad variety of SDC software. The core software will replace the SHELL program to provide a unified framework. It will be designed to support programs doing simulation as well as production reconstruction and analysis tasks.

Early in FY1993, as the core software system is being implemented, design will begin on a new simulation program, intended to be fully operational in the middle of the decade, so that it can be used extensively before turnon of the detector. This program will be used to test reconstruction software for the final detector and to sharpen physics analysis strategies for the first data.

In the short term, we will continue to develop event reconstruction codes and maintain them within the SDCSIM system until the new formal software system is developed. As soon as the new system becomes available, we will reuse codes and/or algorithms from the current system as much as possible. It is likely that we will use the current SDCSIM as a test bench for components of the new system during its development phase.

10.7. Organization

Since computing is a distributed enterprise used by all of the SDC, its management will need to be intrinsically matrixed with that of other subsystems. The computing group must coordinate both the development of a software system that performs data reconstruction and physics analysis and the assembly of hardware at many places that will run the software. The effort must be managed in such a way that many people can contribute to the software development. The computing group must also ensure that computing is available and usable by SDC physicists when they need it.

10.7.1. SDC computing organization

We outline here an organization for computing that will meet our needs. We note that it is not the only way to accomplish the task and probably not, in detail, what finally will be used.

The SDC software will be built around a kernel developed by an SSCL based group most of whom are computer scientists and programmers. Later in the SDC construction and operation, there will be other functions that require a local computing group. Since there will always be a concentration of effort on computing at SSCL, we expect that the computing manager should also be located at SSCL. The major elements of the computing organization we propose are as follows:

- Computing Manager. The manager supervises the SDC computing organization at SSCL, coordinates the efforts of the distributed SDC computing group, and has primary responsibility for ensuring that the SDC computing system is designed and built to carry out the needed tasks on time.
- Area Coordinators. The computing task will be divided into departments such as (a) core software, (b) detector software, (c) networking and communications, (d) hardware, and (e) regional centers, each of which will have coordinators reporting to the manager.
- Detector Software Coordinators. Software for the reconstruction and analysis of each detector subsystem (tracking, calorimeters, muons, etc.) will be developed by groups largely from each subsystem area. Each of these subsystem software groups will have a coordinator reporting to the detector software coordinator.

In addition to the individuals above, a number of committees will be needed to provide coordination and review. A suggested committee structure follows.

- Computing Technical Committee. The technical committee meets regularly with the computing manager to plan the SDC computing system and assess progress on its development. It is composed of the managers listed above and other members selected for their expertise and knowledge in different aspects of computing and SDC.
- Software Committee. This committee is composed of representatives of the core software group and the detector subsystem software groups. It recommends and reviews the overall SDC software architecture and provides coordination among the software groups. It approves standards for software development appropriate to the various parts of SDC software.
- Production Software Review Board. Production reconstruction, being (ideally) a one-time step in the SDC physics chain, has critical requirements on the reliability of the software. It must both carry out the reconstruction correctly and be robust in handling unexpected data without crashing. The Review Board will be composed of computing experts and representatives from each subsystem area. It will supervise and review testing of software planned for inclusion in production code.
- Computing Review Committee. This high level committee of experts both inside and outside of SDC will be convened periodically to provide an overall assessment of the plans and progress of the SDC computing organization.

A computing section within the SDC department at SSCL should contain the dedicated computing effort. The computing manager would supervise the people in the section, among whom would be the area coordinators, the SDC core software development team, and an SDC computing support group.

In addition to the personnel in the SDC computing section, SDC will make use of SSCL computing personnel when appropriate. This will often be arranged as long term assignments of people whose expertise is needed for a period in the development of SDC computing. In addition, SDC will call on SSCL computing experts to work with SDC for shorter periods to design parts of the system or to resolve problems.

10.8. Technology forecast

Much of the technology embodied in the offline computing system will only be obtainable later in the decade. Forecasting is required to provide estimates of future capabilities and costs. Some specific examples are given below.

The current state of the art communication backbone has a bandwidth of 100 Mbyte/s. Progress in optical fiber technology will enable an order of magnitude improvement by the mid-1990's. The campus data highways that attach to this multi-Gbyte complex backbone will be in the 100 Mbyte/s range, replacing today's emerging 10 Mbyte/s FDDI LAN's.

Today's embryonic client-server and peer-to-peer distributed processing support will mature by the mid 1990's. A secure distributed computing environment will allow co-operative processing across institutional subsystems. Progress in distributed systems will result in a system management server to handle job scheduling, load balancing, network management, and hardware and software configuration control.

The collaboration may use multiple large processor complexes. The future TeraFLOP processing complex will be used for large modeling and simulation tasks. It will likely consist of a set of tightly coupled moderate-to-massively parallel processors, possibly utilizing several processor technologies in parallel. Feeding such a processor engine will be high speed mass storage, which may include a matrix of disk arrays supporting data striping and parallel I/O at speeds of 100 Mbyte/s.

Information storage requirements will be met by distributed data management systems, allowing data to be shared across centers. The local processing complex storage hierarchy will be managed via a mass storage archive processor. The system will migrate the center's data from an expanded (64 Gbytes) data cache, through high speed disk farms with capacities of 200 Gbyte/box, through moderate speed tape and optical juke boxes with capacities in the 10 Tbytes range, to longer term automated tape libraries storing 150-200 Gbyte/tape. Archival optical tape systems will be available for permanent off-line storage, with a range of capacities (100 Gbytes on cartridges to 1 Tbyte on large reels). Data migration will be based on usage and user defined migration parameters.

Efficient data compression will reduce storage and communication costs. Lossless techniques will be offered in both hardware and software products, with generic products yielding ratios of 5:1. Custom algorithms should yield 12:1 or more. Generic lossy algorithms will yield ratios of 40:1 or better.

These technology forecasts are summarized in Table 10-5.

10.9. Schedule and milestones

A bar chart of the computing activities needed to provide hardware and software by detector turnon is shown in Fig. 10-8. Conceptual design will continue during the remainder of FY1992. Final design will begin on the kernel portion of the software by FY1993. The kernel program is needed in usable form by the end of FY1994. Design work will begin then for subsystem codes. The schedule provides for a period at the end of the development cycle during which detector-oriented software will be reviewed and tested thoroughly prior to detector turnon.

Hardware for SDC computing will be acquired as late as possible because of the steady improvement in cost to performance ratio that is expected to continue throughout the decade. Design and testing of samples of devices for the SDC will start in FY1993, after the conceptual software work has been done. Small quantities of computing and storage hardware will begin to be acquired in FY1996. The initial devices will allow realistic testing of planned hardware and provide systems on which to carry out the simulations that will ramp up as analyses are prepared for the start of data taking.

Table 10-5
Technology forecasts.

<hr/>			
Basic CPU technology	900K circuits/chips < 5 ns cycle times \$0.0002/circuit		
Processor architecture	Massively parallel (SIMD, MIMD) TeraFLOP processors. Vector super computers with 4-64 processors. Distributed computing (networked processing.) High performance, specialized co-processors.		
Storage	Robotic tape libraries for Petabyte storage. Automated data migration through storage hierarchy. Striping will support transfer rates of 100 Mbyte/s. 64-Mbyte chips will allow 64 Gbyte memory cache devices.		
	Capacity	Data rate	Cost
Disk - Magnetic	20 Gbyte/box	12-15 Mbyte/s	\$0.05-0.07/Mbyte
Disk - Optical	10 Tbyte/jukebox	7-10 Mbyte/s	\$0.04-0.09/Mbyte (no controller)
Tape - Magnetic			
Helical 8mm	5-10 Gbyte	5 Mbyte/s	
Helical 19mm	25-165 Gbyte	15-40 Mbyte/s	
Linear 12.7mm	10-20 Gbyte	10-12 Mbyte/s	
Tape - WORM Optical			
12.7 mm	100-150 Gbyte	12 Mbyte/s	
35 mm	1 Tbyte	12 Mbyte/s	
Communications		Data rate	Error rate
Base Optical Fiber		4 Gbyte/s	$10^{-14}/s$
System Interconnect		4 Gbyte/s	
Wide Area Network		1 Gbyte/s	
Metro Area Network		1 Gbyte/s	
Local Area Network		0.1 Gbyte/s	
Bus width			
64 bit		0.8 Gbyte/s	
256 bit		3 Gbyte/s	
<hr/>			

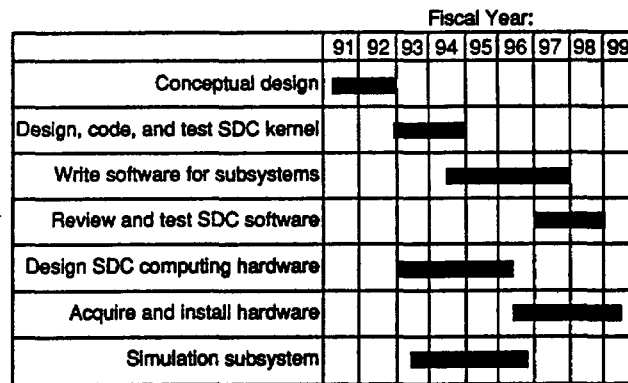


FIG. 10-8. Bar chart of schedule for major SDC computing activities

We are in the process of refining the schedule for hardware acquisition and software production. The present SDC computing hardware and software milestones are shown in Table 10-6.

10.10. Manpower and cost

The costs of offline computing hardware and software are not included in the total SDC project cost (Chapter 15). The cost of providing offline computing for SDC is estimated in this section. For the major hardware elements (processing and storage), a survey has been made of the costs of currently available solutions in each category. A per-year performance improvement factor was then applied. A conservative value reflecting experience of the last few years was chosen in each cost category. The factors applied to correct expected performance at constant real cost are 1.4/year for CPU performance, 1.2/year for networks or switches, and 1.1/year for storage capacity. Hardware costs estimated in this way are shown in Table 10-7. Purchased software associated with the hardware, such as operating systems and compilers, have been included with the hardware costs.

In addition to hardware, manpower costs will be incurred primarily at the SSCL. The manpower is needed principally for software development, but some will also be needed for hardware installation and commissioning. Manpower requirements have been estimated by comparing the tasks involved with similar projects from recent high energy physics detectors. In particular, the experience of CDF and D0 has been used extensively. The software development costs are mostly associated with the development of the kernel software, which is planned to be carried out by software professionals resident at SSCL. The estimate assumes that only 20% of the software for detector systems will be written at SSCL, with the rest produced by physicists, including students and post-docs. The estimate of manpower required is shown in Table 10-8.

Table 10-6
Milestones for SDC computing hardware and software.

Milestone		Date
Fix baseline design		January 1993
Baseline cost estimate		March 1993
Final design		January 1996
Acquire production system	1%	January 1997
	10%	January 1998
	100%	June 1999
Acquire analysis system	10%	January 1998
	100%	June 1999
Identify software development systems:		
Kernel		October 1992
Detector Code		October 1993
Core system:		
Requirements analysis		June 1992
Functional specification		October 1992
Complete design		October 1993
Complete coding Release 1		October 1994
Complete testing Release 1		October 1995
Physics/Detector systems:		
Requirements analysis		May 1994
Functional specification		October 1994
Complete design		October 1995
Complete coding Release 1		October 1997
Complete testing Release 1		October 1998
Simulation:		
Requirements analysis		August 1992
Functional specification		January 1993
Complete design	Rel.1	January 1994
	Rel.2	October 1995
Complete coding	Rel.1	January 1995
	Rel.2	October 1996
Complete testing	Rel.1	January 1996
	Rel.2	October 1997

Table 10-7
SDC computing costs.

Cost Element	FY92 M\$
Production storage system:	
Tertiary storage system	3.70
Communications link to online	0.25
Production computing farm (10 ⁵ SSCUPS)	6.00
Production disk system(s)	1.50
Software	2.00
 Data distribution system	 1.00
 Express line system	 1.00
 Simulation system (10 ⁵ SSCUPS)	 6.00
 Local analysis system	 2.20
 External Networking	 0.25
 Total	 23.90

Table 10-8
SSCL manpower required for SDC computing.

Task	FTE-yr.
SDC Computing Support group 4(1992) · 10(1999)	25
Hardware	
Requirement analysis	3
System design	6
System modeling	2
Procurement	2
Installation	5
Testing	5
Software	
Core system (production and analysis)	
Requirements analysis	5
Design	10
Coding	5
Testing	5
Reviews	2
Documentation	3
Physics/Detector Systems	20
(additional 80 FTE-yr provided by the SDC)	
Simulation software	10
(additional 40 FTE-yr provided by the SDC)	
	118

References:

1. William Atwood *et al.*, "The Reason Project," *Computing for High Luminosity and High Intensity Facilities*, AIP Conference Proceedings **209**, 320(1990).
2. R.J. Ennis and D. Miller, "SDC Computing Model and Operational Concepts," SDC Note SDC-92-229 (1992).
3. C.T. Day, K. Liao, and T.T. Song, "RESQme Studies for SDC Regional Computing Centers," SDC Note SDC-92-230 (1992).
4. SDC Computing Working Group, IBM Federal Systems, *Solenoidal Detector Collaboration Computing Section, Architectural Studies and Analysis*.
5. Kuntal Rawal, "A Macro Facility for X, Object Specific Event Simulation," *The X Resource Issue 1, Proceedings of the 6th Annual X Technical Conference*, O'Reilly & Associates, Inc., Boston, 1992, ISBN 0-937175-96-X.
6. G. Kubena, "Software Engineering Methods and Life Cycle Options," SDC Note SDC-92-231 (1992).

11. Safety

11.1. Introduction

11.1.1. Overview

The requirements for hazard identification, effects analysis, and mitigation to acceptable levels are spelled out in DoE Order 5481.1B, "Safety Analysis and Review System". As a first step toward implementation of this mandate, the SDC has established a Safety Working Group with membership mainly drawn from the engineering design staffs of the major subsystems. The goal of this group is to initiate safety hazard analysis activities during the conceptual design phase of the detector, to facilitate early hazard identification, elimination or control, and to ensure that ES&H considerations are included in tradeoff studies, design alternate evaluations and cost optimizations.

The purpose of this chapter is to summarize the results of the preliminary hazard analyses carried out by the SDC Safety Working Group. This work has concentrated on identifying and analyzing the principal areas of risk inherent to the physics goals and choices of technology selected by the SDC for inclusion in this Technical Design Report. Interfaces between the detector and its utilities and services were defined. Preliminary hazard analyses were also carried out on the significant areas of risk associated with these support systems. The more routine safety issues associated with industrial and research occupancy were not considered at this time unless required to understand detector issues. The Safety Working Group has produced a Draft Conceptual Safety Analysis Report (CSAR) [1].

The CSAR is planned as a working document to be updated and refined in parallel with the detector subsystems design development. As safety analyses broaden in scope and evolve from conceptual engineering judgements to quantitative risk evaluations, the safety documentation will progress from the CSAR to a Preliminary Safety Assessment Report (PSAR), the first of the formal ES&H submissions stipulated by DoE Order 5481.1B. Final designs will subsequently be accompanied by the required Safety Assessment Report (SAR).

11.1.2. Methodology and scope

A uniform process was adopted for conducting hazard analyses and risk assessments of SDC technical system designs and operations. The process follows closely the guidelines in DoE Order 5481.B to facilitate an evolutionary path towards the required documentation, the PSAR and the SAR. Issues of personnel risk received the highest priority. Risks of equipment damage were also investigated in those cases where consequences of a mishap could be critical to the program. Details of the hazard analysis methodology can be found in Ref. 2.

The first step in the process was to conduct a detailed review of technical designs for a given subsystem to identify hazards presented by the design choices or implied operational conditions for purpose of their ultimate prevention, elimination, or control. Hazards which could cause death, injury, occupational illness, or damage to facilities, systems, equipment or the environment were the focus of this step. Hazards not routinely encountered by the public were also considered.

Subsequently, a hazard effects analysis was carried out by the subsystem engineering design groups. The purpose was to assess the probability for mishap and the severity of the consequences prior to risk reduction measures being applied. At this stage, such assessments were not strictly numerical; engineering judgement was applied according to well established definitions for mishap severity and occurrence probability. The consequences could occur in any of four categories: personnel injury/illness, equipment loss, program interruption, and environmental effects. Impact was categorized in a four-step severity range from negligible to catastrophic. World-wide experience with similar, smaller scale colliding beams detectors has provided considerable guidance for these assessments. Experience from the four experiments at LEP and two each at the *SppS* and the Tevatron aided our safety analyses.

The probability and severity assessments were then combined to obtain an overall risk assessment. A simple matrix of probability versus severity was used to map out zones of high, medium, and low risk.

The final step was selection of adequate risk reduction measures for identified high and medium risk technical systems or operational conditions. The guidelines provided to the subsystem engineering teams for resolving identified hazards established an order of preference:

- design for minimum risk, eliminate the hazard,
- design safety systems and devices for hazard control,
- provide detection and warning devices and
- develop safety procedures and training.

Risk reduction measures were then designed for each subsystem and the frequency-severity matrix was revised to establish the abated assessment of risk.

In the following sections, a brief description is given of the more significant hazards presented by the technical designs described along with details of mitigation strategies. The documentation in the CSAR provides considerably more technical detail and includes a full set of hazard analysis forms. These analysis forms serve to document the identified hazards and effects, the initial risk assessments, the recommended mitigation actions, and the abated risk assessments in a simple and uniform manner.

11.2. Tracking

11.2.1. Overview

The charged particle tracking system is situated in a cylindrical volume bounded in radius at 1.75 m by the superconducting coil, in half-length at 4.3 m by the endcap calorimeters, and on the inside by the beam pipe. The basic performance requirements specify measurements of charged particle tracks with extremely high precision while presenting only a minimal amount of material to minimize interactions. The tracking technologies must also be very reliable, since disassembly of the enclosing calorimeter and muon detection systems to effect repair is a major enterprise that is necessarily infrequent.

A small 8 cm crack between the ends of the barrel calorimeter and the endcap calorimeters has been designed to accommodate both supports and services for the tracking detectors. This disk-shaped space also contains supports for the superconducting coil and some calibration services for the calorimeters. The crack will be filled with electrical cables, fiber optic cables, gas lines, cooling lines, power leads, and support members.

11.2.2. Inner tracker

The inner tracker is made of silicon strip detectors inside a radius of 50 cm and half-length of 2.5 m. These are configured as concentric cylinders in the central region and as layers of disks perpendicular to the beam in the forward and backward regions in a way that minimizes the amount of material traversed by particles originating at the interaction point. All of the front-end electronics for the approximately two million individual strip detector elements are located on the ends of the strip detectors with connections made using wire bonds.

The silicon tracker is described in detail in Chapter 3. Of principal concern is the cooling system, which uses evaporating butane to remove heat from the internal electronics.

The cooling system functions as a thermosiphon with the supply liquid and the return vapor piped to an external condenser station for heat extraction. Shutting down power to the electronics results in a return of all liquid to the condenser leaving only vapor throughout the rest of the system. The inventory is less than eight liters of liquid and is primarily determined by the diameter of the liquid supply line and the distance from the tracker to the condenser. Engineering design information regarding the transfer line and the condenser are included in the CSAR hazard analysis.

The probability for and the severity of cooling vapor leaks will be reduced to a minimum by using good engineering practice and complying with industry standards for vessels and associated piping. Details of this aspect of hazard control are presented in the CSAR. In addition, a set of mitigation measures has been required of the design. Automatic and redundant venting will be provided to limit differential overpressure and underpressure conditions in the containment vessel to 0.1 to 0.2 atmospheres. For practical purposes, a design over pressure of one atmosphere will be used in sizing the containment wall. Venting will be engineered to be free of freeze-up malfunctions under foreseeable fault conditions, and the makeup gas for underpressure excursions will be nitrogen. Startup and shutdown procedures have been designed which use an intermediate nitrogen gas step to eliminate the possibility of mixing butane with air. A secondary containment vessel has been designed for the concentric liquid supply and vapor return lines between the condenser and the silicon tracking vessel. The containment space outside of the butane vessels will be actively purged with room temperature nitrogen.

Inside the tracking volume, performance requirements on the amount of material do not allow introduction of a secondary containment vessel around the silicon tracking vessel. Equivalently, the entire outer tracker space is used for leak control. A steady flow of room temperature nitrogen gas is required throughout this space to maintain temperature stability. Using nitrogen for this purpose and sealing off the beam pipe openings through the endcap calorimeters also provides a confining inert environment around the silicon containment vessel. The hydrocarbon and oxygen content of the exhaust gas from this outer tracker space and from the secondary containment space around the transfer lines will be continuously monitored at the 100 parts per million level. Process monitoring of these and other key parameters and operating conditions is planned. Severe off-normal alarms would trigger automatic shutdown of the silicon tracker electronics and initiation of butane removal and nitrogen replacement in the cooling system.

11.2.3. Outer tracking

Two options are currently under consideration for the outer tracker system. This system is situated in the space between the silicon tracker and the inner boundaries of the solenoid and endcap calorimeters. The baseline design specifies gaseous detectors with straw drift tube modules in the barrel region and gas microstrip detectors in the intermediate region. The second option is based on scintillating fibers deployed in an all-barrel configuration which does not require a separate intermediate tracking system. In both cases, the safety and design provisions for electronics described in Section 11.6 are applicable, especially in the case of the gaseous tracking option where all front-end electronics are located within the tracking volume.

Straw drift tube/gas microstrip option

The modular straw tube design utilizes five superlayers of straw modules mounted on graphite epoxy composite cylinders. These concentric cylinders provide a stable base on which to mount the straw modules. The cylinders are themselves supported from the calorimeter using brackets extending into the barrel-endcap calorimeter cracks. The utility interfaces required consist of drift gas lines (presently a mixture of 80% CF₄ and 20% isobutane), high voltage supplies for the sensing elements, low voltage and coolant services for the front-end, trigger and readout electronics, plus a continuous flow of room temperature nitrogen gas through the volume for temperature stability. Electronics connections, drift gas manifolds, and cooling gas manifolds for the straws are located at the extreme ends of the modules; there are no center connections. A total of 720 straw modules are used in the design of the tracker, which produce an estimated 6 kW of waste heat.

The Gas Microstrip intermediate tracker is composed of three discs per end mounted perpendicular to the beam line at the ends of the cylinders of the outer tracker. These detectors are miniaturized drift chambers making use of micro-electrodes placed on glass substrates to establish drift and gas amplification field structures. Front-end and readout electronics are mounted on the back of each of the approximately 10 cm by 10 cm detector units. Utility interfaces to each such detector include three high voltages for the electrodes, drift gas supply and return low voltages for the electronics, and cooling services.

Designs of the detectors and mechanical supports for the gas microstrip system have not yet progressed to a level beyond the conceptual phase. The power dissipation, type of gas, and structural materials have not yet been clearly defined. Xenon or CF₄ based gases mixed with hydrocarbons are the current options being considered. The support structure will most likely be a graphite epoxy composite. Present estimates indicate that the gas microstrip tracker will consume 15 kW of electronics power per end.

The principal hazard issues for both gas microstrip detectors and modular straw detectors have been assumed to be the same for the CSAR analysis. These include combustible gas leaks, loss of cooling accidents, and electrical faults. The risk of fire underlies all of these mishaps. The mitigating actions and precautions described in Section 11.6 for both high voltage and low voltage electronics will apply to these detectors to mitigate personnel and equipment risks to a safe level.

The outer tracker requires a stable thermal environment in order to maintain the alignment accuracy required of the sensing elements while simultaneously minimizing the amount of material in the tracking volume. This means that low-Z materials are used in the structural design, and that temperature gradients induced by the silicon tracker and solenoid surfaces as well as by the high voltage and low voltage waste heat from the straws and gas microstrips must be controlled for alignment stability.

A gaseous cooling system is presently planned to remove electronics waste heat. A continuous flow of room temperature gas distributed uniformly over the cross section of the tracking volume will be used to help maintain thermal stability. As described in Section 11.2.2, nitrogen gas has been selected for this stabilization so as to simultaneously inert the volume and dilute any leaks. The flow rate will be selected to ensure that the exhaust gas will be nonflammable for foreseeable faults in the drift gas systems. The oxygen and flammable gas content of the purged gas volume will be continuously monitored, and high level alarms will be connected to automatically shutdown electronics and high voltage supplies and initiate flammable gas removal/nitrogen replacement.

Design of the gas systems for the straw and gas microstrip tracking detectors will be guided by code requirements, SSCL standards, and good engineering practice to mitigate against leaks. Protection against sudden changes in pressure due to stuck valves, rapid atmospheric pressure variations or human error will be included. Startup and shutdown procedures will be designed which use an intermediate nitrogen gas step to eliminate the possibility of forming air-drift gas mixtures in the system.

A loss of coolant accident or any reduction in cooling efficiency, regardless of origin, would immediately result in overheating of the electronics inside the tracking volume. A multiply redundant safety system is planned to mitigate this possibility. Interlocks will be used to effect a shutdown of all power sources in the event of a severe off-normal alarm condition from coolant flow transducers, valve position monitors, or thermal sensors immersed in the coolant streams or directly in contact with the electronics.

Scintillating fiber option

The scintillating fiber outer tracker design uses six superlayers of fiber ribbons mounted on accurately fabricated graphite epoxy cylinders. These cylinders will provide a stable base on which to mount the fiber ribbons with high precision. As in the gaseous tracker, the scintillating fiber tracker requires a very stable thermal environment in order to maintain the required alignment accuracy. The only utility interface required in the tracking volume is a steady flow of room temperature nitrogen gas - to maintain thermal stability against the silicon tracker and solenoid surfaces, and to provide an inert environment around the silicon butane-cooling containment vessel.

There are no electronics or other active elements in the scintillating fiber tracking volume. Instead, clear wave-guide fibers are used to pipe the scintillation light signals to photodetectors located in the access space around the periphery of the barrel calorimeter. Visible Light Photon Counters (VLPC's) are used as photodetectors. These solid state devices operate at a temperature of 7-9 K. Liquid helium services are thus required for the 64 pool-boiling, vapor-insulated VLPC canisters used in the design. Since the endcap calorimeters can be rolled back by slightly over a meter while the tracking systems remain in place, all VLPC canisters will be supported off the barrel calorimeter to avoid uncabling problems.

The volume of liquid helium in each cryostats is small, of order one liter, so that an oxygen deficiency hazard due to helium boiloff is extremely unlikely even in a catastrophic release. Nevertheless, a full ODH monitoring and alarm system will be installed and connected to automatically trigger the ventilation system's purge function in which the flow is increased by a factor of two using 100% outside air. An accident involving the helium supply transfer line(s) could result in uncontrolled release at a rate of about 50 liquid liters per hour for a few minutes before automatic shutdown of the helium supply can activate. Mechanical protection is planned to reduce the risk of induced damage to the transfer line.

Maintenance activities envisaged include removal of cold VLPC modules and their replacement with warm spares. The hazards to personnel, frostbite and tissue freezing, will be mitigated through the use of protective wear, safe work procedures and practices, training, and design for maintainability. Included in the design criteria for the VLPC cannisters are requirements related directly to the safety of startup and maintenance operations.

11.3. Superconducting solenoid

The superconducting coil is a thin cylinder located just inside the barrel calorimeter; it provides the 2 T solenoidal magnetic field for the central tracking detectors. A chimney-like utility tube connects the coil cryostat to a service port outside the muon system directly above the north end of the solenoid. This service port is the interface between the superconducting coil and the external cryogenic systems and the conventional support services. The chimney runs vertically upward from the coil and connects to the bottom of the service port. Superconducting power leads, diagnostic and sensor connections, cryogen lines, and vacuum and venting services are contained inside. The power supply, control system, and refrigerator are located in the cryogenic system utility building on the surface above the underground hall.

The aluminum cryostat consists of an outer vacuum shell, an inner vacuum shell, and an annular bulkhead on each end. All of the components are installed inside this vacuum vessel, which is supported from the barrel calorimeter. The conductor for the single-layer coil consists of cabled niobium-titanium and copper (Nb-Ti/Cu) superconducting filaments centered in an aluminum stabilizer. Kapton or glass tape and epoxy is used for the turn-to-turn electrical insulation. The coil is bonded to the inner surface of a solid high-strength aluminum support cylinder with a structural adhesive, and cooling is provided by an aluminum liquid helium cooling tube welded to the outer surface of the cylinder. Cold gas return helium cools the aluminum radiation heat shield. Multilayer insulation is used in all of the the vacuum spaces.

11.3.1. Coil

The main electrical hazards to personnel will be mitigated by use of standard insulation, physical barrier, interlock, and operating procedure measures. Also, the power supply will be designed, installed, and tested in accordance with applicable codes and standards and with good engineering practice. Magnetic field hazards to personnel include both health and work hazards. The high field region of the system cannot be accessed when the coil is energized due to the enormous magnetic pressure acting on the calorimeter endcaps; thus magnetic health hazards are not present.

Rigid policies, work rules, and administrative procedures are planned regarding activities inside the tracking volume to preclude direct induced damage by workers and indirect induced damage due to magnetic fields acting on tools or other items left behind after maintenance. Work hazards could also arise due to magnetic forces acting on tools and equipment being used in the fringe field regions. If work is necessary in an area of sizable fringe field, mitigating actions would include training, a field strength survey, and safe work procedures, practices, and equipment.

The coil operates at 8000 amps and has a stored energy in excess of 100 MJoules. In the event of a quench, the coil energy must be dissipated as rapidly as possible while avoiding completely the severely damaging conditions of excessive coil temperatures, excessive coil voltages, or open circuiting. Quench protection relies on redundant detection and automatic switchover from the power supply to the dump resistor. Continuous monitoring of the integrity of the dump resistor is foreseen. The integrity and

calibration of the detection system is automatically verified with each charge-up or discharge of the coil. The primary causes for unwanted quenches, insufficient cooling, small sudden mechanical motions, and local regions of very high field, are mitigated *ab initio* by the design criteria. Finally, a prototype coil will be built to validate the engineering design and verify the predictions for the forces acting on the coil prior to initiating fabrication.

Materials control measures are imperative to eliminate the possibility of components or materials that do not comply with design specifications. Very high stresses from magnetic pressure and thermal contractions are handled by specifying special high-strength materials. General mitigation actions planned include verifying physical and chemical properties by obtaining material test reports or certificates of compliance, using material identification marking systems, and by direct in-situ testing.

11.3.2. Cryogenics

The primary perceived hazard is the spill of a cryogenic liquid in a confined area leading to its vaporization and displacement of oxygen. Without appropriate mitigation, a cryogenic spill onto a warm surface results in rapid vaporization followed by resettlement of heavy gases with little or no mixing with surrounding air. In a confined area, this leads to an asphyxiation hazard unless measures are taken to dilute or remove the gas. General mitigating actions universally applied for such hazards will be used; these include designing to limit the quantity and rate of accidental releases, providing for emergency purge ventilation, designing and fabricating to ASME codes, providing oxygen monitors and ODH alarm devices, and requiring personnel training and protection devices.

Secondary cryogenic hazards to personnel, frostbite, tissue freezing and hypothermia, are primarily an issue of prevention by adequate design. Details of the design rules and codes used for various parts of the coil, cryostat, and cryogenics system to reduce the probability for a release to a minimum are presented in the CSAR. In special circumstances where more direct exposure could be part of the operating procedure, preventative measures would include training, use of protective wear, and work procedures designed for safe work practices.

Loss of insulating vacuum failures can have consequences that range from a minor increase in the heat load to uncontrolled vaporization leading to severe overpressure conditions. Reliefs for limiting pressures to design values are required for each of the separate cryogenic vessels in the system. These are sized to handle the flow from a maximal foreseeable accident, usually an abrupt loss of vacuum. The design strategy calls for separating the vacuum requirements into as many independent isolated spaces as is practical and for the use of redundant pumping stations. Extensive monitoring of the vacuum system is planned to alarm on early trends and to automatically trigger mitigating actions for serious off-normal occurrences.

11.4. Calorimeter system

The SDC calorimeters are sampling devices. The overwhelming bulk of material in these massive structures is inert metal absorber sufficient to range out and absorb the energy of particles produced in SSC beam-beam collisions. Layers of relatively low density particle detectors are alternated with the thick metal layers to sample the showers developed in the dense metal absorber and provide a measurement of the energy of incident particles. A high degree of both lateral and longitudinal segmentation is required. Structural safety issues for the calorimeters are discussed in Section 11.8.

11.4.1. Central and endcap calorimeters

Plastic scintillator has been chosen for the detection layers in both the barrel and endcap calorimeters. Each scintillator piece or tile has a wave-shifting plastic fiber embedded in it to collect the scintillation light and re-emit it at longer wavelength. A clear wave-guide plastic or silica fiber is used to connect the scintillator tile/fiber detectors to photomultiplier tube readouts located on the outer calorimeter surfaces. Local crates for front-end, trigger, and readout electronics are also mounted on the calorimeter back surface. High voltage and low voltage hazards in the calorimeter access space will be minimized through compliance with the mitigation measures and standards described in Section 11.6.

Both the plastic scintillator tiles and the plastic fibers are combustible hydrocarbons. Approximately 55 tons of plastic material is used in the central calorimetry system. For practical reasons, a modular approach to fabrication and assembly was selected. Each separate module must be a light-tight unit to avoid interference from room light. The light-tight container for a module is formed from sheet metal. Consequently, the fire risk is greatly reduced as the combustible plastic is protected and will not be accessible as fuel in the early stages of a fire incident. Moreover, the 60 to 1 ratio of inert metal thermal mass to combustible plastic fuel, and the absence of ignition sources, further serves to limit the fire risk.

Significant magnetic fields will permeate parts of the calorimeter as the flux return for the superconducting coil is effected through the central calorimetry absorber steel. In particular, several hundred tons of magnetic force are exerted on the endcap calorimeters, pulling them towards the barrel. Multiply redundant prevention measures will be used to completely eliminate the possibility of energizing the coil with an endcap partly removed; all such measures involve inhibition of the coil power supply. Independent mechanical and optical proximity interlocks are planned to sense proper endcap positioning. Power supply lockout and tagout requirements will govern any movements of the endcap calorimeters. Operational readiness signoffs will be required prior to power supply turn-on. Finally, the possibility of sensing or interlocking the fasteners which secure the endcaps to the barrel is under investigation.

11.4.2. Forward calorimeter

Forward calorimeter assemblies are located about 12.5 meters from the interaction point. These cone-shaped calorimeters are comparatively modest in size, weighing approximately 150 tons per end. Radiation doses up to 100 megarads per year are possible immediately adjacent to the beam pipe. As a result, the consequences of radioactivation of calorimeter materials as well as the physics performance requirements must be considered in choosing detector technologies for this region.

Instead of alternating layers of metal plates and sampling readout detectors, the forward calorimetry designs all use long, thin fiber or spaghetti-like detection elements that penetrate the absorber from front to back. The two candidate detector technologies are discussed briefly below.

Liquid scintillator option

The conceptual design for the liquid scintillator option consists of a steel absorber matrix interspersed with 5 mm diameter quartz tubes. Liquid scintillator would be circulated through the quartz capillaries at a rate sufficient to achieve the necessary level of survivability. A complete replacement of the liquid inventory would be required whenever the activation levels of the fluid approached the threshold for classification as mixed low level radioactive waste.

Clear fibers inserted in the liquid filled tubes are used to bring the scintillation light signals to photodetector readouts. The flexibility of these fibers allows for a low hazard approach to placement of the readouts and associated services. In particular, the readouts can be located entirely above the containment vessel, eliminating the need for penetrations below the free standing liquid level and their associated leak exposure.

If this option is selected, an evaluation of the total combustible load and risk of ignition will be undertaken to permit design of adequate mitigation and control measures. Where risk of leakage exists, collection gutters, pans, and sumps will be engineered and installed. Vapor and liquid accumulations in

the collection points will be monitored. The choice of the scintillator chemical composition would take into account the production of long lived isotopes, as well as performance requirements. Activation studies and measurements would be required to determine isotope production rates and detector useful lifetime as well as shielding and handling precautions for maintenance operations.

Pressurized argon gas option

The conceptual design for the high pressure gas ionization option consists of small, 10 cm by 10 cm cross section, modules utilizing a steel absorber matrix interspersed with 3 meter long pressurized argon tubes. Inside the approximately 7 mm diameter tubes, ionization signals would be collected on wire or rod electrodes running down the tube axis. The design pressure is 100 atmospheres, and approximately 800 modules would be used per end. The ionization gas, argon, would be replaced at a frequency dictated by performance survivability; no problems of low level radioactive waste disposal are anticipated.

If this option is selected, the individual detector modules would be required to be designed in accordance with applicable pressure vessel codes and manufactured as code-stamped pressure vessels. To set the scale, a module is comparable in volume and stored energy to a standard K-size pressurized gas cylinder. An evaluation of asphyxiation hazards would also be undertaken to design indicated mitigation measures. Activation studies and measurements would be required to determine isotope production rates and detector useful lifetime as well as shielding and handling precautions for maintenance operations.

11.5. Muon system

The central, intermediate, and forward muon detection and triggering systems comprise the outermost layers of the SDC detector. The muon detection system consists of wire chamber measurement layers, magnetized iron absorber layers, and scintillator trigger layers. This Section is concerned with the muon detection systems and the magnetizing coils. A description of the structural safety issues relevant to the magnetized iron absorbers can be found in Section 11.8.

11.5.1. Drift chambers

The basic sensing element in the muon tracking system design is a round aluminum tubular drift chamber. Outside diameters of 45 mm, 60 mm, and 93 mm are used in different regions of the system and lengths vary from 1 to 9 meters. Each tube is a complete and independent drift chamber and will operate slightly above atmospheric pressure. Individual tubes are assembled into larger mechanical modules appropriate to the geometry of each layer. Such a modular assembly of tubes has associated with it a gas supply and return manifold, a high voltage and preamp electronics assembly, and a set of alignment fiducial devices.

Since the size of the muon wire chamber system is conveniently measured in acres (presently about 12 acres), the most significant safety issue concerns the flammability and toxicity of the gas mixture. The drift tubes have been engineered to meet the SDC performance specifications using a mixture of Argon and Carbon Dioxide as the drift gas. This eliminates the possibility of drift gas fire and toxicity hazards.

Other potential muon chamber hazard such as structural integrity of the chamber assembly supports, high voltage distribution and low voltage distribution for the local electronics are discussed in Sections 11.8 and 11.6 respectively. The remaining risk areas that are associated with the unusual size and scope of the chamber system include a potential oxygen deficiency and confined space hazard, an interior assembly structural failure possibility, and the requirement for working at significant heights for maintenance and installation.

The primary SDC response to the potential for oxygen deficiency in normal and possibly confined spaces associated with the muon chamber system is to provide for an extensive oxygen deficiency sensing and warning system (see Section 11.10). Situation awareness, as provided by such a system both prior to entry and during access to the chamber areas shall be designed to prevent personnel exposure to ODH environments. Training, medical screening, and personal monitoring will be among the principal hazard

mitigation requirements enforced for access to the muon chamber areas interior to the detector. Additional information on access policies with respect to designated confined spaces is found in Section 11.11.

Other mitigation measures under study concentrate on reducing the probability for a major gas leak. The most significant of these is establishing a maximum working pressure parameter for each subsection of the system and installing safety relief and flow restriction devices appropriate to that pressure limit. Such relief systems shall be capable of safely venting the maximum possible flow in a fault condition and shall be located so as to guarantee their function. The muon chamber gas is recirculated allowing for direct leak detection by comparison of the supply and return flow rates, by monitoring the oxygen contamination level in the return gas, and by monitoring the consumption of make-up gas. Also, maintenance scenarios are to be carried out by trained personnel following well established procedures. These procedures are designed to mitigate hazardous operations, and promote personnel and equipment safety.

11.5.2. Scintillation counters

Plastic scintillator with photomultiplier tube readout is used as a fast response detection layer in the muon system to provide the prompt timing needed to "tag" or identify the particular beam crossing responsible for the "hits" in the drift chambers. The safety of the photomultiplier readouts is addressed in Section 11.6.1 on high voltage hazards. Although a decision has not yet been made on the exact material involved, all plastic scintillators are pure hydrocarbons. There is only one significant hazard in the usage of tens of tons of these materials, namely hydrocarbons are combustibles.

SDC fire protection strategies, as described in Section 11.7, are predicated on the universal installation of incipient fire detection equipment. The very early warning thus provided forms the basis for mitigation strategies and control methods. In the case of the muon system plastic scintillator, incipient detection will be supplemented by a prevention measure. The scintillator will be enclosed in metal containers so as to deny availability of this material as fuel in the early stages of a fire incident. A non-trivial side benefit of containerizing the scintillator concerns the installation and maintenance activities around the detector and the subsequent lack of a need for "do not step here" signs on otherwise exposed scintillator surface. The container is assumed to remain intact until external heat sources have raised temperatures of the entire assembly to the melting point of the plastic.

Under active study, is the issue of manual, local application, and automatic fire extinguishment systems for the scintillator layer. Because of incipient detection, a staged response capability is highly desirable to permit suppression actions appropriate to the severity of a fire incident. The requirements for such an extinguishment system are that it shall not present an unacceptable direct threat to life, and that fire control shall not, on average, cause more damage than a probable fire incident itself. Solutions to this problem will have many applications throughout SDC.

11.5.3. Magnetizing coils

The steel absorbers in the barrel and forward muon detection systems are magnetized to allow a measurement of track momentum that is independent of the central tracking measurement. The barrel and each layer of the forward absorbers are toroidally magnetized to 1.8 T using a set of water cooled copper coils. The current will be approximately 4500 amps, and the corresponding stored energies are 18 MJ and 5 MJ for the barrel and forward toroids respectively. Because of geometrical constraints, the barrel coils will be made in two pieces, and assembly of a field joint is necessary on installation.

The principal hazards to personnel, high voltages and high currents, can be mitigated by using standard insulation, physical barrier, power supply interlock and operating procedure measures. For the risk to equipment, the two main mitigation measures are extensive testing and quality control during fabrication and an electrical and thermal coil safety system. Techniques planned for the former include high potential insulation tests and capacitor discharge electrical tests, and such tests would be repeated in-situ for final acceptance and for pre-run checkouts. The safety system is based on sensors monitoring the operating conditions of the coil and power supply system, and it is directly and redundantly interlocked to the power supply mains contactor. Included in the monitoring are cooling water conditions, current and voltage

values, coil and joint temperatures, power supply operating conditions, and ground fault current. Since failures of the safety system could imply a loss of protection, the safety system design will use redundancy and supervision along with periodic testing to mitigate against alarms and actions failure.

11.6. Electronics

Electronics refers specifically to the three major signal processing functions necessary to acquire event data from the detector: front-end electronics, trigger electronics, and data acquisition systems. In addition, the high voltage systems needed for operation of drift chambers and photomultiplier tubes are included under electronics.

Conventional electrical systems for power, lighting, convenience outlets, etc. are not discussed here as the ES&H issues are well covered by compliance with OSHA requirements and the National Electrical Code. For electronic systems, there is no analogous code or handbook of good engineering practice. Despite the fact that electronics systems are characterized by low voltage DC power systems, typically 15 volts and lower, there is still a significant fire and thermal damage hazard because of the possible high current capability of the supplies.

11.6.1. Overcurrent protection

World HEP experience has shown that lack of adequate overcurrent protection and temperature monitoring in low-voltage high-current distribution networks and electronics systems are the most frequent causative factors in fires in experimental facilities. The same is true for high-tech facilities in general, from research laboratories to telephone switches. The most significant vulnerability occurs where power sources are mounted physically separate from the loads introducing a low voltage distribution network between the supplies and the loads. This separation, however, is usually the optimum configuration for installation, maintenance, and cooling, and is also the one recommended by safety analyses of operational and maintainability issues.

SDC has elected to implement a pro-active overcurrent protection policy to ensure that designers and users of low-voltage high-current electronic systems shall take all reasonable steps to assure safe operation under foreseeable fault conditions. The policy calls for mandatory design criteria with operational readiness reviews and inspections prior to initiating operations. A set of design and implementation safety criteria have been drafted and are described in the CSAR. The underlying concept is that all current carrying power conductors shall be protected in accordance with their current-carrying capacity. The accompanying design review and operational readiness inspection policies will be developed in conjunction with other similar SDC controls to maintain a globally coherent oversight situation.

11.6.2. High voltage

High voltage, greater than 1 KVDC, is required for operation of wire chambers and photomultipliers in the tracking, muon, and calorimetry systems. To achieve a set of practical design requirements and mitigation measures, usage of high voltage has been divided into two classifications based on life safety; those which can be directly lethal as measured by the threshold for onset of ventricular fibrillation and those which cannot. The electrical parameters related to shock severity are the current and duration of a fault condition, and the stored energy. The duration time limit parameter is especially relevant to installations with active fault detection and shutdown (or trip) capability. The suggested limits for these parameters to distinguish between the two classes of hazard have been taken from the data on human effects contained in International Electrical Commission publication 479-1, Effects of Current Passing Through the Human Body.

Current Limit	10 mA
Duration Limit	20 msec
Stored Energy Limit	10 joules

For installations characterized by parameter values less than the above limits, there is still the possibility for secondary accidents, *e.g.* falls caused by electrical shock, and mitigation measures are designed to minimize such exposures. In these installations, the design shall ensure that no live parts are accessible when voltage is on. Any work on energized systems shall be controlled by a hazardous activity permit procedure, and conspicuous labeling shall be used. Finally, special care shall be taken in the selection of external high voltage connectors to ensure that connection operations engage the neutral or return conductor first, and that an unplugged energized cable cannot spark to an external object or person.

For installations with any one parameter in excess of the above limits, additional measures are required. During operation, total inaccessibility of the high voltage shall be assured by physical barriers and/or interlocks. Access for repairs, modifications, etc. shall be controlled by a mandatory lockout and tagout procedure. In the case of stored energy hazards, remotely actuated discharge systems shall be used to make the installation safe prior to access. Each element of the installation shall be clearly labeled with adequate warnings.

For purposes of increased high voltage safety, photomultiplier systems shall consider designs which develop the high voltage directly in the tube base. In this way, high capacity (bulk) high voltage power sources and extensive high voltage distribution and interconnect systems can be avoided. Safety of the corresponding bulk low voltage supply and distribution system would be governed by the provisions of the overcurrent protection policy described in Section 11.6.1.

11.6.3. Rack protection

The standard NEMA relay rack is designed as a stand-alone cabinet for installation of chassis mounted electronics and crate systems. Before modification by users, the rack is fully contained with three sides, a top and a bottom; the front is closed by electronics assemblies and blank panels. Cable penetrations and cooling requirements are the usual reasons for modification. Such rack units then only present hazards in accordance with the character of the installed equipment; the most significant of these is thermal damage and fire exposures from high power density systems. Removal and control of waste heat is a primary performance and safety design concern.

SDC planning calls for removal of waste heat from rack-mounted equipment by means of embedded air-water heat exchangers. It is anticipated that each rack in the data acquisition system could contain up to 12 KWatts of electronics and the corresponding low-voltage power supplies. This air-water choice is not only the most effective in controlling the operating temperatures of the equipment, but is also the most efficient when considering the alternate of adding these loads to the HVAC plant. A properly designed rack cooling system could contribute to the total air conditioning requirement by operating slightly below ambient temperature.

Although not explicitly called out in code specifications or DoE Orders, SDC has identified such high-power rack systems as presenting significant risk. Rack protection, designed to mitigate against the inherent fire and thermal damage hazards of 12 KWatt rack power densities, is included in the engineering design requirements. The mitigation strategy is based on an extension of existing requirements for process monitoring: detection of off-normal voltages, currents, and temperatures. Equipping each rack with an independent power distribution panel, as would be normal for a 12 KWatt load, allows for the inclusion of a shunt-trip main breaker, interlocked to sensors monitoring system integrity. Detection of a minor fault would result in an off-normal condition alarm; detection of a major fault would result in complete shutdown of the rack by means of the shunt-trip main breaker. Among the process conditions presently planned to be monitored for operational integrity, the following have been identified as being of potential significance for automatic shutdown:

- smoke detection in the rack air stream;
- temperatures above and below each crate;
- temperatures of power supplies;
- cooling water flow through heat exchangers;

- cooling water supply and return temperatures;
- cooling water leaks;
- overcurrent conditions in power supplies; and
- overvoltage conditions in power supplies;

It is to be emphasized that the rack protection system proposed is autonomous; severe local off-normal conditions only power down the affected rack. This is in contrast to the more global protection systems designed in accordance with accepted methodology for protection of electronics areas as a whole. Power shutdowns triggered by alarms of a global nature, as described elsewhere in this chapter, will be effected through the power mains for the entire complex, not on an individual rack basis. An example is loss of room air conditioning where the off-normal condition is overtemperature in the entire area and the mitigating action is dropping all power to the same geographical area. Rack protection can be viewed as an extension of such global systems to the conditions within each separate rack.

Examples of rack protection systems can be found at major HEP laboratories presently using FASTBUS, VME, and MULTIBUS electronics. These include CERN, Fermilab, SLAC, TRIUMF, and BNL among others. The SDC expects to design and implement a rack protection system taking full advantage of this operating experience. Among the currently active areas of inquiry and research is a search for a local-application rack based extinguishment system; this was triggered by the announced ban on new halon installations. Substitutes are not yet available. Another important area of inquiry is the use of sub-atmosphere water cooling circuits which offer lower probability against damaging leaks. The final rack protection system, as specifically designed for SDC, will be submitted for technical peer review to a special engineering group convened for the purpose.

11.7. Fire protection

11.7.1. Overview

In fire protection orders and cited codes, DoE has placed emphasis on four principal areas; life safety, program continuity, property protection, and releases to the environment. The size and scope of the SDC detector presents some very important challenges to fire protection engineering. By its very mission, the detector is not capable of being partitioned into physically separate fire risk zones of lower value. Also, the deep underground location compounds problems of providing for emergency egress and smoke ejection. The situation is much like that found in many high-tech, high-value facilities where the probability for a fire incident is intrinsically very low, but the consequences could be catastrophic in the absence of protection and mitigation measures. At this point in the development of an integrated and comprehensive fire protection plan for the SDC detector, the priority is on developing those concepts, albeit qualitatively, which directly affect technical designs for the detector systems or pertain to Title I design requirements for the underground facilities.

The most significant combustible loading in the detector is plastic scintillator in the calorimeters and muon trigger layers. Design measures taken to reduce the exposure and make the plastic inaccessible as fuel in the early stages of a fire incident are described in Sections 11.4 and 11.5. In both the calorimeter and the muon designs, the scintillator will be encased in metal canisters preventing direct exposure.

The next most significant combustible loading is the cable plant. The most important of the mitigation measures planned is to use cable types that have the lowest toxic and corrosive smoke rating and flammability index consistent with electrical performance requirements. Essentially all fire accidents in HEP and similar high-tech facilities have had the same experience; the secondary damage resulting from the generation of corrosive smoke significantly exceeded the direct thermal damage of the fire itself.

Finally, modest quantities of flammable gas are used in the tracking system situated in the center of the detector surrounding the beryllium beam pipe. Even though a fire incident here would be relatively insignificant itself, strong prevention and mitigations measures are called for because of the potential

exposure to the beryllium beam pipe and cable plant. These measures are described in Section 11.2, and include continuous monitoring for leaks with automatic shutdowns on alarm and inerting the tracking volume with a steady flow of nitrogen gas.

Fire protection is being planned for the SDC complex according to the activities and contents of each separate area or zone. In the vast majority of situations, provisions of the Life Safety and NFPA codes will be followed for each of these local occupancies. In the few cases where the codes are not relevant or impractical, mainly concerning the detector itself, a principal of equivalency will guide the design. Hazard-specific protection and prevention systems will be designed to achieve the same or superior levels of protection as indicated in the codes. In this section, the conceptual plans for fire protection of the detector in the collision hall are described.

11.7.2. Detection

Incipient detection is planned for all interior spaces within the detector and the overhead space immediately above it. In the majority of HEP fire incidents, investigations have shown that technical problems or fault conditions existed for anywhere from 1/2 to 24 hours before the accident. An aspiration, sample-draw smoke detection system is envisaged for the different internal layers with modularity appropriate to the natural openings, maintenance accessways, and combustible loadings. In addition, a high-sensitivity spot detection system will be deployed that is based on an evolved gas signature. Carbon monoxide and hydrocarbon sensors that are sensitive in the few parts-per-million range will be used to provide the early warning. The two systems provide redundancy and work together to provide very early detection. The spot detection system also provides location information unavailable from an aspiration system alone.

Multitiered alarming will be used to take maximal advantage of the early warning from incipient detection. At the lowest level, the alarm condition is sent only to the local operations center for diagnosis of the condition. At the middle level, the off-normal condition is sent electronically to relevant operational service centers, and investigative response is mandatory. An extensive system of TV surveillance cameras is planned to allow remote investigation of pre-alarms. At the highest level, a fire alarm is generated and appropriate actions and annunciations are triggered.

11.7.3. Suppression

Given that incipient detection will be installed, the suppression equipment should be designed to allow for a staged response that provides for any incident. The goal is to allow for localized fire control appropriate to the scope of the alarm that would minimize the induced damage of the fire control method. In the past, such systems have used extinguishment equipment that ranged from standard portable carbon dioxide units, to wheeled tanks of halon 1211 with hand hose lines, to water standpipes with hose cabinets, to total flooding halon 1301 systems, to full sprinkler wetting of the collision hall.

With the recent ban on the use of halon in new facilities, this broad range of low damage fire control agents is no longer available. Existing alternatives such as total or local flooding with carbon dioxide, nitrogen, or high expansion foam have a major drawback in that they are directly life threatening. Lockout and tagout of the suppression system would be required during personnel access and maintenance periods. Much of the advantage of incipient detection would be lost waiting for the access control system to validate complete evacuation. The SDC, along with the SSCL Physics Research Division, has formed a fire protection advisory panel whose membership includes fire protection engineers from both National Laboratories and high-tech industries. The development of fire protection plans, including staged suppression systems, will be done in consultation with this panel.

11.7.4. Prevention and loss control

A series of significant prevention and loss minimization measures have been identified during the CSAR hazard analysis process. These are listed and briefly described below.

Automatic Shutdowns— A fire alarm will automatically shut down all electrical power to the detector and initiate flammable gas removal and nitrogen gas replacement.

Smoke Ejection— A fire alarm will cause the HVAC equipment to switch into an active smoke ejection mode.

Response Procedures and Training— On-site support personnel will be trained to execute investigation and mitigation procedures for both pre-alarms and fire alarms.

Overcurrent Protection Policy— Mandatory engineering and inspection standards are being developed for overcurrent protection of low-voltage high-current electronics systems.

Rack Protection— A rack protection system has been specified to automatically shut down individual racks based on local sensor information including smoke detection.

Flammable Gas Safety System— All areas containing flammable gas will be monitored and fitted with multitiered alarming. Automatic shutdown of all potential ignition sources and changeover from flammable gas to nitrogen gas follows detection

Tracking Volume Inerting System— The tracking volume will be inerted with a steady flow of nitrogen gas.

Process Monitoring— A very extensive system of process monitoring is planned for the slow controls. Those off-normal conditions that are directly relevant to fire risks will be separately alarmed as potential fire risk warnings requiring mandatory investigation.

Housekeeping Policy— For many reasons, including fire prevention, a strict housekeeping policy is envisaged.

11.8. Structures

Structural integrity of the SDC detector and its components is not only essential to safe and reliable operation during the life cycle of the experiment, it is also vital to performance. In order to satisfy alignment stability requirements, several of the mechanical support structures are required to meet deflection and stability specifications well in excess of those required by applicable structural codes. The policy adopted by SDC requires that supports, components, lifting fixtures, and assembly fixtures for the detector shall be designed and engineered in compliance with the guidance listed in applicable codes for allowable stresses and safety factors to ensure overall structural integrity. Moreover, the mechanical design criteria shall be reviewed in accordance with the appropriate DoE directives for hazard analysis and review. Finally, the design criteria will be reviewed for alignment stability performance. A partial list of applicable codes follows:

AISC,	"Manual of Steel Construction"
ANSI A58.1,	"Minimum Design Loads for Buildings and Other Structures"
AWS D1.1,	"Structural Welding Code"
DoE 6340.1A,	"General Design Criteria"
ANSI/ASME B30.20	"Below-the-Hook Lifting Devices"

In general, the adequacy of more conventional detector structure designs will be reviewed as part of the hazard analysis process for each individual subsystem. Examples are toroid coil support brackets, beam pipe supports, and wire chamber supports. The size and scope of the SDC detector encompasses very large scale structural components and supports that are also complex, procedure oriented, and structurally critical. These unusual structures have been identified and described in the Structures Section of the CSAR

to allow for a unified hazard analysis and mitigation strategy. Major mechanical and structural components of the detector that are included in this separate analysis are:

- Barrel Toroid and Support
- Forward Toroids and Supports
- Barrel Calorimeter and Support
- End Cap Calorimeters and Supports
- Forward Calorimeters and Supports
- Calorimeter Transfer Bridges
- Muon Module Support Structures

Preparation of design criteria, calculations and documentation is considered essential for all structural systems associated with the SDC detector, especially those listed above. The adequacy of SDC design criteria shall be the domain of the safety review process, now initiated with issuance of the draft CSAR. The accuracy of SDC design calculations shall be the domain of peer reviews through the SDC technical review process. Such design criteria and calculations shall be prepared, reviewed, and approved prior to commencement of fabrication or installation activities. In addition, conventional safety reviews are required to assure that non-structural safety concerns associated with assembly, detector operations, and detector maintenance have been addressed by the designs.

The technique for review of calculations is expected to vary depending on the details of the individual component design. These will range from hand calculations for relatively simple configurations to execution of a fully independent finite element analysis for a complex critical joint or bolting pattern. The working criterion for the latter, at present, is that such numerical cross checks shall deal with potential input errors, modeling errors, or math errors by using a completely different finite element analysis code or by changing both the boundary conditions and the modeling on the same code package.

Finally, the design criteria shall not be limited to static situations. A very important additional requirement follows from a determination of dynamic situations and deflections and stresses set up by the movement of extremely heavy loads during assembly and maintenance or resulting from probable failures such as that of one jack in a multijack system. Design criteria shall be developed such that supports and structures are engineered to accomodate such anticipated local dynamic deformations, including that of the floor itself, in addition to the static operational loadings.

11.9. Radioactive materials

Radioactive materials will be found in two areas of the SDC detector and the collision hall. First, the central calorimetry uses a system of permanently installed radioactive sources for calibration purposes. Each of the separate sources is attached to the end of a flexible wire which is normally coiled up inside of a lead storage container and shield mounted on the outside of the calorimeters. During calibration operations, a source is pushed through one of numerous small metal capillary tubes, which serve to guide it to precisely located calibration positions deep inside a calorimeter stack. Because of the frequency with which calibration data is taken, the source mover system will be automated for remote operation.

The SSCL/PRD source inventory and control policy now under development will govern the acquisition, packaging, personnel protection measures, and dosimetry requirements for usage of the calibration sources. Design of the lead storage container and shield shall use criteria limiting radiation fields in potential personnel access or work areas to the specified standards. Shipment of the sources or source mover assemblies is governed by these and appropriate over-the-road regulations.

The second set of radioactive materials results from the steady exposure of detector components to secondary particles produced in beam-beam collisions. Significant activation levels are predicted for the calorimeters and the accelerator components close to the beam pipe, as well as parts of the beam pipe itself. At the radius of the beam pipe in the forward/backward calorimetry regions, ionizing doses of 100 MRad

per year are estimated at design luminosity. Dose levels decrease dramatically with distance from the beam pipe, thus there is a naturally defined zone of radioactivation.

All SDC and accelerator equipment in and adjacent to the activation zone will satisfy certain criteria. Such criteria shall preclude generation of mixed radioactive waste, and shall anticipate and mitigate *ab initio* potential problems of ultimate decontamination and decommissioning of the facility.

Designing for maintainability is of crucial importance to keep worker exposures as low as reasonably achievable by minimizing the frequency and duration of maintenance tasks and by maximizing the worker separation distance and self-shielding potential of the structures. Portable and semi-portable shielding assemblies will be required for temporary protection in high radiation areas as any permanent shield would itself become activated. Use of real-time dosimetry is planned to ensure compliance with SSCL/PRD ES&H Manual listed policies and practices for radiation workers.

SSCL/PRD policies and procedures will govern the labeling, handling, removal, and storage or disposal of radioactivated materials. Mandatory training requirements are foreseen as an integral part of such exposure prevention measures.

11.10. Atmosphere management system

Due to the varied nature and amount of gases and cryogens utilized by the SDC detector, awareness of the quality of the atmosphere both inside and outside of the detector is crucial. The SDC detector and collision hall will incorporate an integrated Atmospheric Management System (AMS) providing for constant monitoring and assessment of atmospheric variables and initiating appropriate actions and responses in the event that an abnormality is detected.

11.10.1. Atmosphere management requirements

Detector subsystems utilize both cryogenic liquids and several room temperature gases as part of their system design. The last few superconducting quadrupole magnets of the low beta focussing system will also be located in the collision hall. The two accelerator tunnel openings at either end of the hall will be sealed off to isolate the collision hall volume. Preliminary requirements for gas and cryogen usage are listed below.

Tracking Systems	Nitrogen inerting gas Nitrogen cooling gas CF ₄ -Isobutane drift gas Butane cooling liquid and gas
Solenoid Systems	Helium liquid and gas Nitrogen liquid and gas
Calorimeter Systems	Argon gas
Muon Systems	Argon/CO ₂ drift gas
Accelerator Systems	Helium liquid and gas Nitrogen liquid and gas

The AMS will be required to be able to detect off-normal conditions resulting from mishaps in any of the above systems, provide information useful for locating the source, and provide continuous analog concentration readings to allow design of automatic and manual intervention responses appropriate to the scale of an incident.

11.10.2. System features

One of the key functions of the integrated AMS for SDC is to provide continuous monitoring of the breathability of the atmosphere in the various maintenance access spaces inside and under the detector and throughout the collision hall. To optimize the effectiveness of alarms and automatic mitigation actions and to provide useful local situation awareness reporting, the AMS sensors will be organized into independent protection zones. A preliminary estimate is that eight zones will be needed: the detector pit, the tracking volume, the calorimeter access space, the muon access space, the north forward region, the south forward region, the north collision hall zone, and the south collision hall zone.

In addition to monitoring breathability, the system will monitor combustible gas levels to guard against undetected leaks. Multitiered alarming is planned because of the ability to achieve reliable detection at concentrations significantly below the lower explosive limit. A low level reading will be brought to the attention of the local operations center crew for diagnosis and investigation. A pre-alarm condition will be annunciated electronically at the relevant operational service centers, and investigative response is mandatory. Extensive process monitoring information is available from all of flammable gas systems to aid in the diagnosis. At the highest level, provisionally set at 25% of the lower explosive limit, a full flammable gas alarm is initiated and all appropriate shutdown and mitigation actions will be triggered automatically. In particular, all power is shutdown to preclude ignition sources; the AMS and other life-safety services, however, will use explosion proof techniques and remain energized.

During normal operation, the AMS will be monitoring the integrity of the inerting system for the tracking volume. As described in Section 11.2, a continuous flow of nitrogen is used in the tracking volume for several purposes including dilution and inertion of possible flammable gas leaks. The presence of oxygen is then an off-normal condition requiring investigation and repair. Multitiered alarming is planned to provide early warning and allow time to correct the problem at an early stage. At 10% oxygen, flammable mixtures are possible; a high level alarm set at this point would be used to trigger shutdown and mitigation actions.

As an additional level of sensitivity to off-normal gas or cryogenic fluid conditions, the AMS will be interfaced to the slow controls system. Important gas and cryogen parameter values are then available for monitoring and alarm purposes. Among the variables considered at this point for such inclusion are the pressures and flows in all of the supply and return lines.

11.11. Confined spaces

Due to complexity and number of electronic and mechanical parts associated with the SDC detector, it is expected that some maintenance will be required on a routine basis. In an attempt to provide an effective and safe maintenance environment, design engineers have clustered components requiring maintenance into two main internal detector areas. Access to the majority of photo/electrical transducers, initial signal processing equipment and preamplifiers is provided in the calorimeter and muon access spaces. The former is located between the outer boundary of the barrel and end cap calorimeters and the inner boundary of the muon system, and the latter is an annular space between the outermost two layers of muon detector modules. Access to the detector mechanical support structure, a small portion of the muon electronics, and the alignment jacking system is located under the detector in the deep pit area. As a result of difficult access and equipment congestion and because of the actual and potential hazards present in these areas, both the deep pit and the inner detector access spaces have been classified as confined spaces.

11.11.1. Confined space definition

The SSCL ES&H organization and ANSI Z-117A define a confined space as any area whose primary function is something other than human occupancy, has restricted ingress and egress, and contains actual or potential hazards. Traditionally, a confined space is categorized as being either, "Permit Required" or "Non-permit Required" depending on the level of risk.

A permit-required confined space is an area which, after analysis and evaluation, is found to contain actual or potential hazards to the extent that 1) prior authorization is required for entry and 2) the area requires constant monitoring during access. During detector operations, the inner-detector maintenance areas are permit-required confined spaces.

A non-permit-required confined space is an area which meets the general definition of a confined space but which, after analysis and evaluation, is found to have less severe hazards or has had hazards abated by engineering controls. The detector deep pit area is a non-permit required confined space.

11.11.2. Inner detector access spaces

The inner detector access spaces are needed for maintenance of electronics systems, crates and power supplies, photomultiplier tubes, calibration systems, alignment systems, and cabling. The superconducting coil chimney, a conduit for the coil utilities, and a variety of gas and water system services pass through these areas. Hazards within these spaces can be categorized into three groups; oxygen deficiency hazards, electrical hazards, and mechanical hazards.

Potential oxygen deficiency hazards are a result of both cryogenic and gaseous conduits existent in the area. Cryogenic helium, nitrogen gas, and detector gases pass into or through the areas. To abate this risk, the areas will be extensively monitored using both oxygen and hydrocarbon gas sensors. An integrated atmospheric management system, Section 10.10, will provide real-time awareness of both the oxygen levels and presence of other gases in the environment. In the event of detection of hazardous atmospheric conditions, personnel will be alerted to perform safety measures for which they have been trained and the area will be evacuated. Automatic shutdown and ventilation enhancement actions will also be triggered to mitigate the hazardous conditions.

These inner detector maintenance areas provide access to electronics crates servicing the tracking, muon, and calorimeter detector systems. This equipment requires both high and low voltage support. All electrical equipment in the area will be designed and installed in accordance with both the National Electrical Code (NEC) and the SDC standards for electronics and high voltage installations described in Section 11.6. These required safeguards will be augmented with procedures and training

An important characteristic of this confined maintenance access space is its physically constrained nature providing both mechanical hazards and limited entry and exit. This peculiarity becomes exaggerated in an emergency situation when immediate exit or the extraction of an incapacitated individual is required. The layouts for the inner maintenance access spaces are being planned with the above in mind. Electronic and mechanical components are being designed for easy fault isolation, removal and replacement. Aisles, ladders, and crawl spaces are required to allow for human access. Personnel access routes to equipment areas, with special emphasis on the lower three barrel octants, are being planned not only for facilitation of maintenance operations, but to allow for the removal of a disabled worker as well.

The access spaces are being designed in conjunction with the detector hall to maintain an emergency exit time of less than five minutes from anywhere inside the detector. Overall, the inner detector access spaces are being designed to provide a safe reliable, predictable environment for qualified personnel. In keeping with laboratory policy, safety measures in this confined space strive not for compliance but for excellence. Emergency lighting, power and communications planned exceed DOE and OSHA requirements.

11.11.3. Detector deep pit area

Access to equipment for the support and alignment of the SDC detector is located in the collision hall deep pit area. While this area does not present the same level of hazard as the interior access spaces, it does meet the fundamental requirements for a confined space and is thus being designed as a non-permit-required confined space.

Continuous monitoring of the deep pit air quality is included in the SDC integrated atmospheric management system to provide real-time situation awareness information. In the event of detection of hazardous atmospheric conditions, personnel will be alerted to perform safety measures for which they have been trained and the area will be evacuated. Emergency ventilation enhancements will be automatically triggered. The mechanical aspects of the deep pit are being designed to allow for human factors, needs and emergency situations. As with the interior access space, emergency lighting, power and communications strive for excellence, exceeding mandated minimum requirements.

11.12. Summary

The technology choices and corresponding specific implementation designs presented in this Technical Design Report have been subjected to a hazard identification and effects analysis. An SDC Safety Working Group was formed for this purpose, the members of which were chosen from engineering design teams for the different detector subsystems. Although such hazard analyses are conceptual at this stage, involving engineering assessments as opposed to numerical evaluations, the goals of early identification of hazards, of assessment of the level of risk and of severity of consequences, and of initiation of risk abatement activities have been achieved.

Subsequent to the identification process, hazards were classified in degree of risk according to probability for occurrence and severity of consequences using a matrix method derived from DoE instructions and guidance for a safety analysis and review system. Where serious potential risks were found to exist, prevention or mitigation measures were identified that would eliminate or abate the risk to acceptable levels. In all of the detector subsystems, it was possible to identify such abatement measures. Details of the analyses and mitigation recommendations where needed and of the risk assessments before and after abatement are contained in the Draft Safety Working Group Report [1].

It is the judgement of the SDC Safety Working Group that the detector designs described in this Technical Design Report can be engineered to achieve adequate levels of safety with respect to personnel injury/illness and the environment. With respect to equipment loss and program interruption risks, the mitigated assessments do not all fall into the lowest risk category due to remote probabilities for mishap coupled to critical or catastrophic consequences. Such identified hazards will be the focus of continued engineering efforts to reduce accident probabilities to the minimum achievable.

References:

1. "Draft Conceptual Safety Analysis Report" (CSAR), April, 1992, SGT-000003,
2. "Hazard Analysis Instructions," Draft Project Management Document on SSC Laboratory Practice D10-000003,

12. Experimental facilities

The SDC detector will be located at interaction region IR8 located on the east side of the collider ring near the town of Palmer. The SDC detector requires a large underground experimental hall with associated surface buildings for assembly, experimental equipment, offices and utilities. These facilities contain the required technical and conventional support systems for the detector—*e.g.*, cryogenics, gases and magnet power. This chapter provides a summary of technical requirements and present concepts for facilities and conventional utilities associated with the SDC experimental program. More detailed requirements, concepts, and schedules may be found in the SDC Experimental Facilities User Requirement (SEFUR) document [1].

12.1. Site layout

The collider and the east experimental areas are shown in Figs. 12-1 and 12-2. The proposed layout of SDC surface facilities is shown in Fig. 12-3. The IR8 site benefits from the close proximity to major transportation routes, municipal services, and utilities.

Many of the subsystems elements are large and heavy. The weights and sizes to be transported to the site may be as large as $10 \times 10 \times 3 \text{ m}^3$ and up to 100 tonnes. Even larger and heavier components will be moved between the surface facilities at IR8 during the construction of the detector. The calorimeter subassembly is roughly $4 \text{ m} \times 10 \text{ m}$ diameter and weighs 1500 tonnes. Special attention will be given to the design and usage of the road system both on and off the site. Access roads to IR8 and interior roads will allow adequate turning radius for transport vehicles.

A rail head is available at Palmer for off-loading detector components or materials which may be brought by rail to the site.

Gas, water, and power utilities associated with the site infrastructure are required for both the construction and operation of the detector. These services are available in close proximity to the site.

12.2. Surface facilities

Surface facilities have been designed to support the assembly, installation, and operation of the detector. The site configuration has evolved from the study of the subsystems assembly, detector installation, and detector operations. A plan view of the site facilities is shown in Fig. 12-3.

Many of the pieces that make up the detector subsystems will be developed and produced in various parts of the U.S. and abroad. Some components may be built elsewhere at the SSCL site. Subsystems will be assembled in the surface facilities at IR8 and then installed in the underground hall. An assembly building with adjacent structures for assembly processes, material storage, and staging will support these efforts.

Table 12-1 summarizes the configuration of the surface facilities.

12.2.1. Assembly building

The major assembly area is in the IR8 Assembly Building (Fig. 12-4). The surface facility requirements are driven by the necessity to perform subsystem assembly and to reduce installation time in the underground hall. Installation is reduced by building large and nearly complete assemblies prior to installation in the underground hall. The assembly area will be designed and equipped to accommodate the subsystems assembly. The location of the assembly building at the IR8 site is such that the heavy calorimeter assemblies can be moved directly to the installation shafts of the underground hall.

The assembly area must accommodate the assembly of more than one subsystem. The major assemblies to be performed in this facility are the tracker systems, muon chambers and calorimeter sections. The layout in Fig. 12-4 shows the initial configuration of the facility when the calorimeter and muon chambers are being assembled. The facility is configured to provide component staging, test areas, storage, shop areas and office space. Though many of the assembly requirements for the subsystems are similar, some

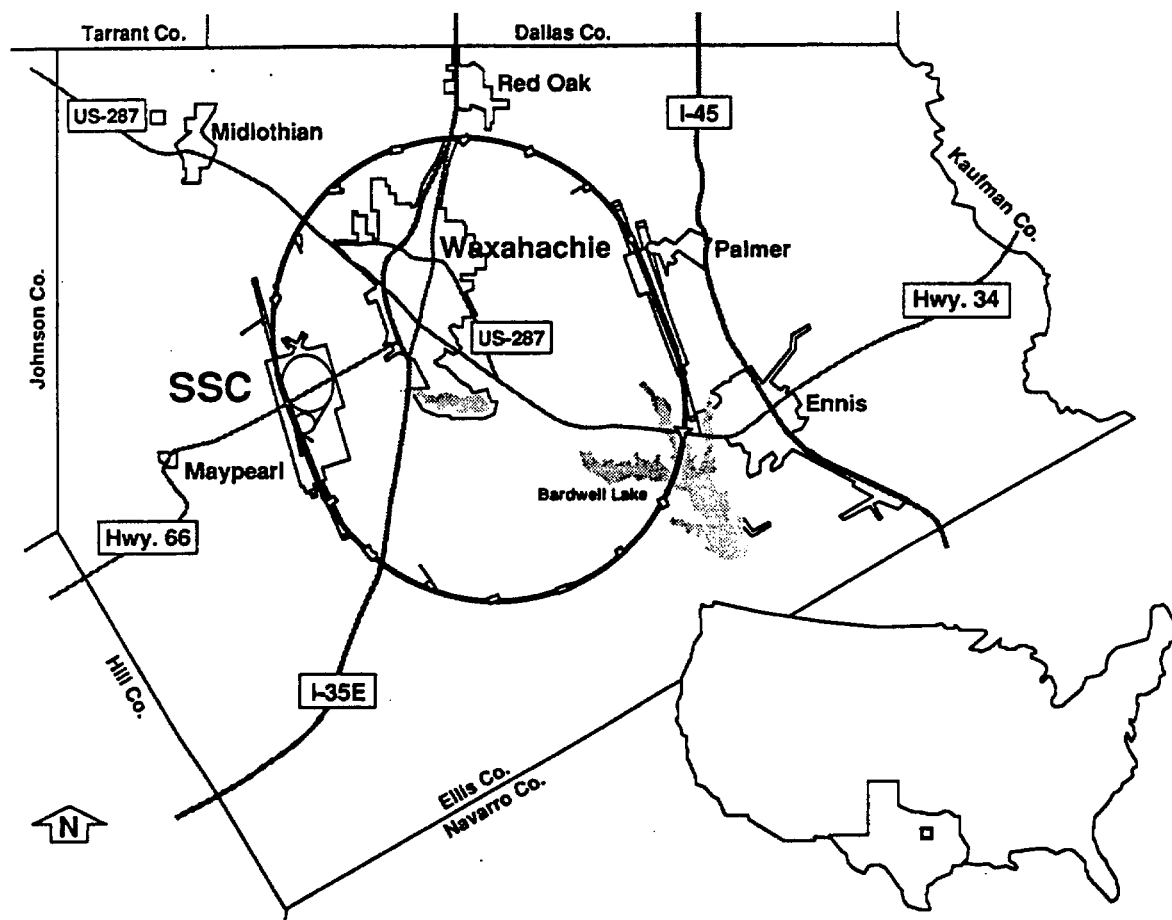


FIG. 12-1. The SSC site.




FIG. 12-2. The IR8 site.




FIG. 12-3. IR8 site plan.

Table 12-1
SDC surface facility requirements summary.

Building Identification	Functional Use	Floor Area (m ²)	Clear Hgt (m)	Floor Loading		Crane Cap (Tonnes)	Env. Req.
				Max. Equip. (Tonnes)	Concentrated (Tonnes)		
IR8 Assembly Bldg	Assembly	3400	12 m	3400	600	1@45/18 2@18/45	HVAC
	Office			(Office)			HVAC
	Storage			40	10		HVAC
Installation Headhouse	Staging/ Access	792	11	2000	500	90/18	Vent
Utility Bldg.	Cryo Systems	1920	10 m	(Heavy Industrial)		None	Vent
Personnel/Equip. Hdhse.	Hall Access	256	4 m	(Light Industrial)		N/A	HVAC
Det. Ops. Building	Ops. & Comp.	2240	5 m	Office/(Computing)		N/A	HVAC
Gas Mixing Bldg.	Detector Gases	149	10 m	(Light Industrial)		None	Vent
Staging Hardstand	Storage	2000	N/A	100 tonnes	50 tonnes	N/A	None

are unique and may require some modifications to the assembly areas to accommodate the differences. A sequence of uses might start with the space configured for muon chamber module storage, then reconfigured to assemble the calorimeter, and then a complex of "clean" areas and alignment facilities installed for tracker assembly.

12.2.2. Detector operations building

The Detector Operations Building houses operational support for the experimental hall and the detector. The control room is located here where the detector is operated and monitored along with the safety systems. The building also contains the off-detector electronics system such as trigger, data acquisition and online computer. This building will provide areas for offices, meetings, computers and light laboratory space for the operating crew.

A system of cable trays are used to bring data cables for the detector and hall electronics into the lower level of this building. These link the detector with the trigger and data acquisition systems, which in turn are controlled by the online computing system in the rooms on the ground floor.




FIG. 12-4. Layout of the IR8 assembly building during muon chamber and calorimeter assembly.

12.2.3. Utility buildings

The Gas Mixing and the Utility Buildings house the utilities for the detector, underground and surface facilities. The Gas Mixing Building, shown in Fig. 12-5, contains gas mixing and monitoring equipment for the detector subsystems.

The Utility Building, shown in Fig. 12-6, will house the mechanical and electrical equipment that services the following:

- Site and underground heating systems (boilers).
- Site chilled water systems.
- Site high temperature, low conductivity water system.
- Site motor generator system (415 Hz).
- Site emergency power system.
- Site and underground energy management system.
- Utilities building power distribution and motor control systems.
- Detector low temperature, low conductivity water system.
- Detector DC power distribution systems.
- Detector cryogenics systems.

Storage tanks, transformers, and air handlers to support the equipment in the building are located adjacent to the building. An emergency power generator will also be located adjacent to the building to provide emergency power for emergency systems, alarm systems, sump pumps and elevators in the underground hall in addition to cryogenics systems and selected systems in the surface facilities.

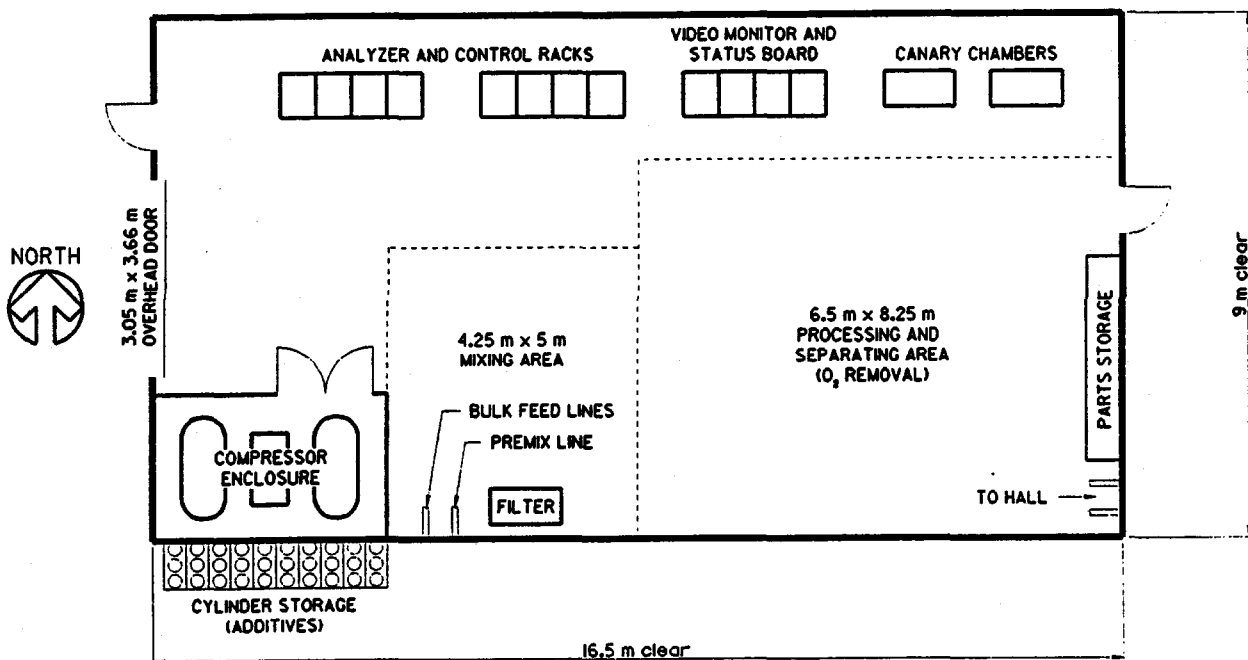


FIG. 12-5. Layout of the gas mixing building.

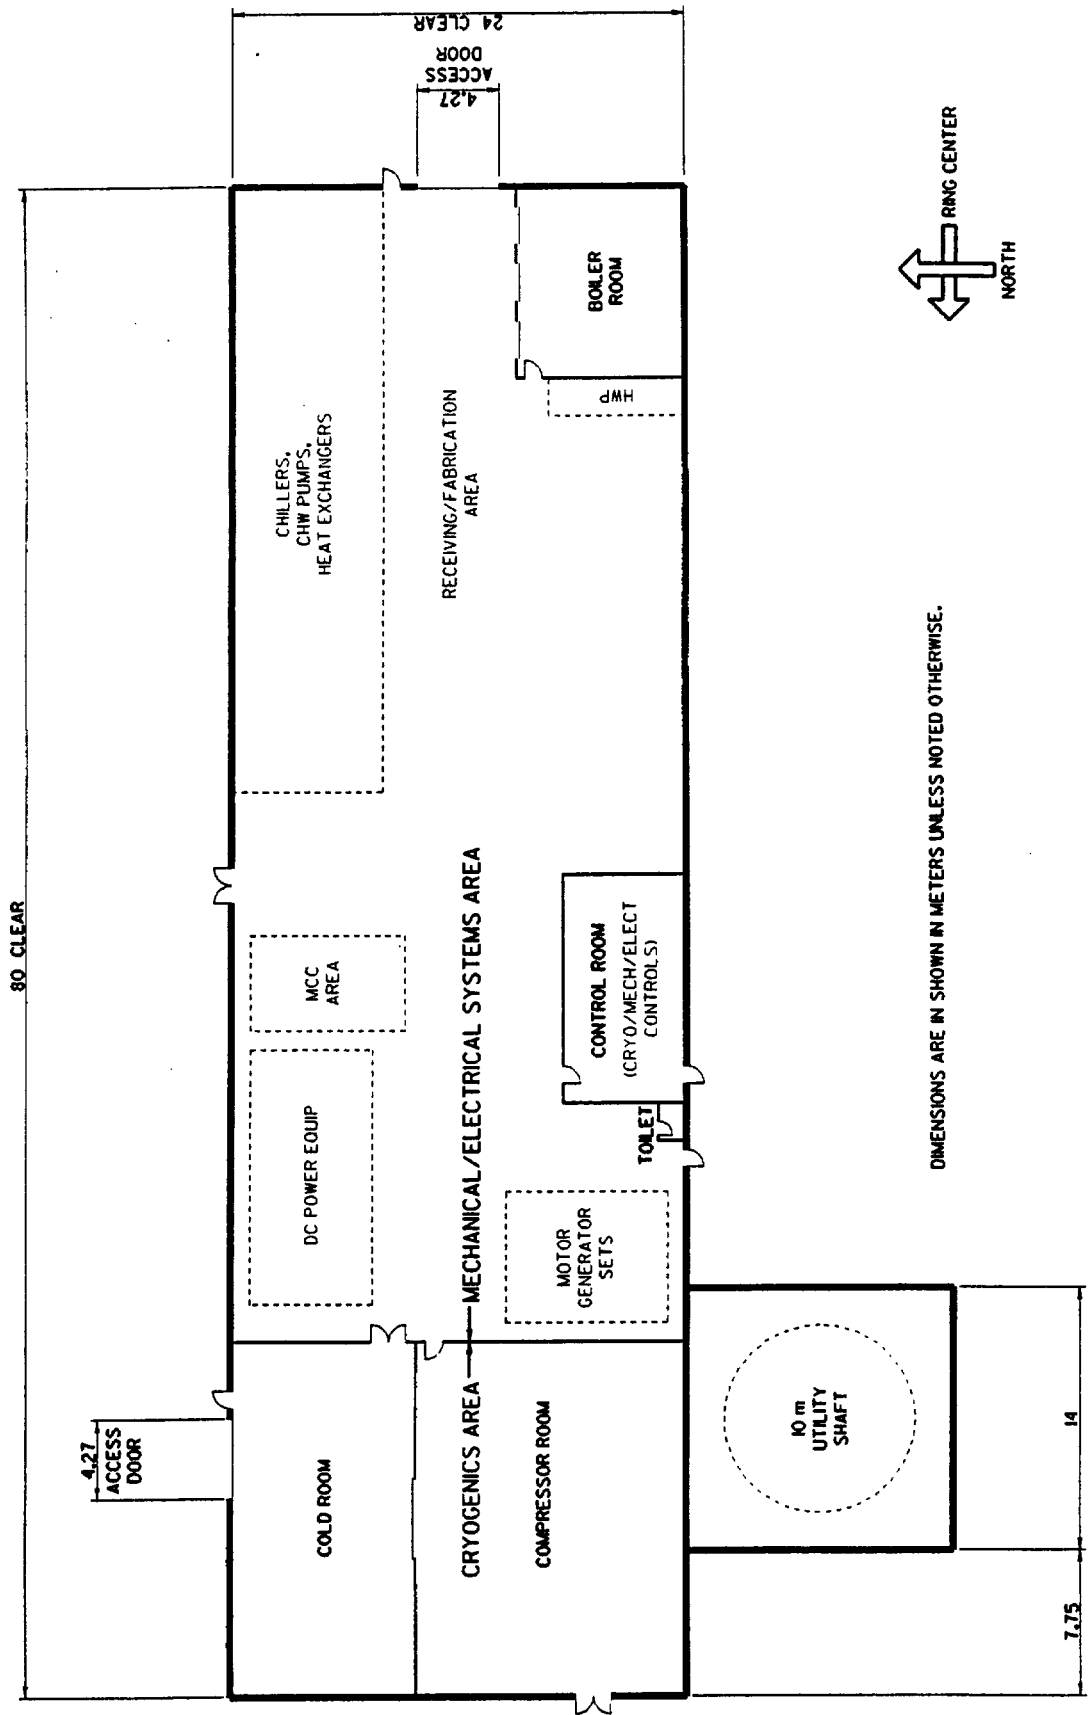


FIG. 12-6. Layout of the utility building.

12.2.4. Installation headhouse and installation shafts

There are two 11×14 m installation shafts connecting the surface to the experimental hall. The north shaft will be covered with a weather tight removable cover.

Over the south shaft, the Installation Headhouse provides a staging area and lifting equipment for installation of detector components into the underground hall. The staging area is sized to allow assembly of the forward muon chambers. A gantry crane spans the headhouse, the areas north and south of the headhouse, and the north installation shaft. This crane will accommodate placing 100 ton pieces up to 10×10 m in size into the underground hall through either shaft. Specialized lifting equipment will be used to lower the calorimeter subassemblies into the underground hall since they can be large objects weighing up to 1500 tonnes.

After detector installation is complete, each installation shaft will be closed with a minimum of 8 m of removable radiation shielding.

12.2.5. Personnel & equipment access building

Equipment and personnel access to the experimental hall during the installation and operation of the detector is through the Personnel and Equipment Access Building. It will be equipped with an elevator and stairwell and will serve as an access point for personnel. The access building also provides personnel access control and radiation shielding during collider operation.

12.2.6. Office facilities

Two permanent office areas are planned; one in the Operations Building and the other in the Assembly Building. The Assembly Building, shown in Fig. 12-4, will provide office space for approximately 50 people early in the assembly and installation period. The Operations Building will house approximately 20 people. The remaining population will be housed in other facilities of the SSCL and in temporary quarters at the IR8 site. After installation is completed, the Assembly Building may be configured for additional office space.

The Assembly Building will be the first structure built on the IR8 site and will establish an early presence for the collaboration and laboratory support personnel on the east campus. Beneficial occupancy of this building is expected early in 1994.

12.2.7. Cranes and lifting equipment

Several types of lifting equipment will be employed at the IR8 site. Lattice boom, mobile hydraulic, and gantry cranes will be used extensively during the installation period of the detector. Specialized lifting and moving equipment will be used to transport large and heavy components within the assembly facility and a special hoisting system will be used in the headhouse for lowering loads in excess of 100 tonnes.

The Assembly Building will be equipped with overhead bridge cranes to facilitate subsystem assemblies. These will be incorporated into the building structure to insure full use of the floor space in the assembly areas.

12.2.8. Temporary staging areas

Areas will be provided for the temporary staging of materials and components for the detector subsystem assemblies. They will also be used for some fabrication tasks and possibly for storing completed modules prior to installation in the underground hall. Two staging areas are planned: a preparation area for the muon toroid steel and a general receiving area. Most staging areas will be constructed of a compacted granular fill but, in cases where condensation may be a problem, the area will be paved and sealed.




FIG. 12-7. SDC underground hall.

12.3. Underground facilities

The underground experimental hall is shown in Fig. 12-7. The detector is situated parallel to the beam line and centered at the interaction point (IP). The dominant factors driving the configuration of the underground facility are the detector configuration, installation plan, maintenance requirements, and radiation shielding. The hall floor is level but not parallel with the beamline which drops by 0.2 m from the north end of the hall to the south end.

The length and the width of the hall are driven by the detector installation and maintenance sequences. The detector is positioned between the two raised operating floors, north and south, in a pit to allow heavy items inside the barrel toroid to be rolled in and out of the detector directly from the operating floors. The south floor is larger than the north floor since a larger portion of the assembly and maintenance operation takes place there. The locations of the two installation shafts with respect to the underground hall floors are shown in Fig. 12-8.

During installation, a small machine shop and offices will be located in the hall.

The personnel and equipment shaft located on the east side of the hall contains a stairs and a 4000 pound elevator. The vertical shafts have been offset from the underground hall to reduce radiation levels at the surface. Access from the stairs and elevator to the underground hall is at the operating floor. Both the staircase and elevator are pressurized for safety reasons. The minimum height for the elevator opening is 2.5 m and the minimum foot print will be no less than 3 by 4 m.

The utility shaft contains all technical and conventional utility systems. This shaft is also configured to reduce radiation levels on the surface.

Each wall of the experimental hall has gangways for access to the detector and mounts for the survey monuments as shown in Figs. 12-9 and 12-10. The gangways are primarily for personnel access. The ones at the hall ends provide access to accelerator components and the accelerator tunnel as well as to emergency egress paths.

Collider components are also located in the hall. Isolation from the hall foundation will be needed to provide for the low beta quadrupole magnets. During long maintenance periods the forward muon toroids may be extracted for access to internal detector pieces. The quadrupoles and their support structure will have to be removed to provide such access.

Two 90 tonne cranes with 18 tonne auxiliary hooks will be used in the experimental hall for assembly and disassembly of the detector. The cranes will be capable of being operated simultaneously (but independently) from separate portable control stations and also be capable of tandem operation from a single control station.

12.4. Services

Gas, water, and power utilities associated with the site infrastructure are required for both the construction and operation of the detector. They are summarized in Table 12-2.

Table 12-2
Summary of infrastructure requirements

Power	15.2 MW
Natural Gas	TBD
Water:	
Potable	34,000 liters/day
Non-Potable	34,000 liters/day
Waste	17,000 liters/day




FIG. 12-8. Plan view of the SDC underground hall.




FIG. 12-9. Isometric view of the SDC underground hall showing gangways and access points.




FIG. 12-10. End view of the SDC underground hall.

12.4.1. Power

Electrical power for the IR8 site will be supplied by three circuits from the East Main Substation. The IR8 substation circuits will be used for a feeder to the Detector Operations Building, a feeder for the Utility Building and a feeder for the balance of the surface facilities. The circuits will have the largest possible separation of electrical buses to provide isolation between the technical and the conventional electrical loads. Manual transfer switchgear at the substation will allow a load transfer should a circuit be taken out of service. Secondary distribution voltages will be 4160V, 480Y/277V, and 208Y/120V, 3 phase, 60 Hz. Table 12-3 summarizes the IR8 power requirements.

A central site emergency power source will be required to provide standby power for all life safety and emergency loads at the IR site. This will be supplied by a diesel powered generator at 480V Y and transfer switching located at the Utility Building.

The clean power circuits for the Operations Building are at 480Y/277V. Clean power is available to electronics within the building and is also routed down the cable shaft to transformers which connect to branch distribution panels for use by the 208Y/120V electronics located in the hall or on the detector. A small uninterruptable power system will supply power for a orderly shutdown of data and control equipment.

Normal power circuits for the Utility Building are also routed down the utility shaft at 480Y/277V for distribution to pumps, welding outlets and other heavy motor loads. This voltage is transformed to 208Y/120V for receptacles, small motors and light welding equipment.

The Utility Building contains the paralleled 415 Hz motor generator sets with an output of 208Y/120V for the detector-mounted electronics. These circuits are routed down the utility shaft and into the hall through isolating transformers. The transformer and high/low secondary outputs are connected to distribution panels located at the hall walls.

The emergency power distribution switchboard circuit is run down into the hall to a branch panel used to provide power for equipment and lighting requiring a backup power source. Switchgear will be mounted outside each building in walk-in type enclosures.

Lighting in the hall uses high bay fixtures with metal halide lamps. Lighting in the shafts utilizes PL-type lamps in wall pack type fixtures. Emergency lighting for the hall and shafts is powered by a battery-backed inverter system in the Utility Building.

Table 12-3
IR8 site power summary

Location	Est. Load KVA	Service Voltage	Remarks
Surface:			
Conventional	2064	480Y, 208Y	
Technical	269	480Y, 208Y	4160 for Cryo
Underground:			
Conventional	2024	480Y, 208Y	
Technical	9772	480Y, 208Y	
Emergency	970	480Y, 208Y	
Total	15099		

12.4.2. Communications

Site utilities will also include signal, control and communication conduits, both between and within buildings. The telephone service will be routed underground in the utility service corridor to the Operations Building and then be distributed from the central switch to the underground hall.

12.4.3. Water systems

The majority of the waste heat rejected to the cooling water systems will be transferred to the atmosphere by means of evaporative cooling systems. These systems will be cooling ponds, cooling towers, or a combination of both. Location and other factors make it necessary to utilize air-cooled sensible heat exchangers to reject some portion of the cooling water system load. The water cooling systems are classified as: cooling pond water, low conductivity water, and chilled water.

Cooling pond water (CPW) is used directly to cool the cryogenic compressor oil coolers and the high temperature low conductivity water, and as condenser water for the water chillers.

Low conductivity water (LCW) is used to cool magnets, magnet power supplies, and the magnet DC power bus, or any application where low conductivity is required for insulating reasons. This water will be deionized and polished to maintain a $10 \text{ M}\Omega/\text{cm}$ resistivity. LCW is required at two temperature levels. The first is the low temperature LCW, used for the muon toroid magnets only (65 to 85°F) and the second is the high temperature LCW used to cool the magnet power supplies and power bus.

Chilled water (CHW) is used for coils in air conditioning units and equipment not requiring cooling water with insulating properties, but requiring cooling water with a supply temperature below what is achievable by evaporative cooling alone. Chilled water is produced by mechanical refrigeration and is supplied at two different temperatures. CHW directly from the chillers is available to air handling units where 45°F is acceptable. Mixed chilled water (MCHW) is produced by a heat exchanger where CHW is the primary cooling medium and MCHW is the secondary water supply. The temperature of the MCHW is kept about 5°F above the dew point (generally, the space dew point is 55°F which requires the MCHW to be no cooler than 60°).

12.4.4. HVAC systems

Air systems are important for detector support, both during operation and when maintenance is needed. Several systems are provided to guarantee that all requirements are met.

The hall ventilation system provides a source of air to assure sufficient oxygen for occupancy and sufficient outside air to dilute any gases escaping from the detector. The hall atmosphere can serve as a heat sink for the small percentage of heat given off by the detector not removed by the detector electronics water cooling systems. For this purpose, air handling units are furnished on the surface which constantly mix 25% outside air with 75% of the return air before it is conditioned. When an emergency such as a fire occurs, these units can double their airflow and use 100% outside air for smoke removal.

The detector ventilation system provides a source of 100% conditioned air for the interior of the detector. During run periods, this air serves as a means of removal for any gases escaping in the detector, and as a heat sink for the small amount of heat not removed by the water cooling systems. During maintenance this air supply serves as a source of fresh air for occupancy. To maintain a low temperature gradient through the detector, the temperature of the supply air is specified as 70°F and is designed to rise no more than 5°F before it is exhausted.

The common vent system is designed to collect nonflammable gases discharged from the detector, vacuum pump exhaust, cryogenic relief valve discharge, mechanical room exhaust and other nonflammable gases.

12.4.5. Gas systems

The Gas Mixing Building will house the gas systems equipment. The mixed gases will be argon and carbon dioxide for the muon chambers and CF_4 and isobutane for the straw tube option. They will be delivered to each detector chamber at the required pressure, flow rate, temperature, purity, and gas constituent ratio according to each systems technical specifications.

The silicon tracker will have a butane cooling system that will reject its heat through a condenser into the MCHW return line. The silicon tracker, and possibly the entire tracking system, will be inerted by gaseous nitrogen.

12.4.6. Cryogenics

Helium refrigerator/liquefier system

The helium refrigerator/liquefier cryogenic system is used to support the superconducting magnet and, possibly, Visible Light Photon Counters. The required refrigeration/liquefaction capacity is approximately 1200 watts and the liquid helium flow rate to the solenoid is 90 g/s.

The cryogenic helium refrigerator/liquefier is housed in two rooms of the Utility Building. An acoustically insulated room contains the helium gas compression system. This system is composed of the compressors, oil removal and purification systems, coolers, and other support equipment for the compressors. Another room contains the refrigerator/liquefier cold box module, air compressors, liquid helium pump box, vacuum pumps, and a regeneration skid. High and low pressure gaseous helium storage tanks and a liquid helium storage dewar are located adjacent to the building.

The liquid helium supply transfer lines are routed from the output of the refrigerator/liquefier cold box to the utility shaft, down the shaft to the utility tunnel, and into underground hall to the detector. The return transfer lines follow the same path as the supply lines. This path length is approximately 130 m. The cryogenic lines are routed in as straight and short a manner as possible to minimize pressure drops and heat leaks. This also provides less complicated fabrication and lower costs. The interface point between the refrigerator/liquefier system and the detector cryogenics system is at the input of the control dewar which is located on the top of the detector.

Liquid nitrogen system

The liquid nitrogen supply system is used to support the thermal shield in the magnet cryostat and the helium refrigerator/liquefier. The supply is from a storage dewar that is located adjacent to the Utility Building. A liquid nitrogen subcooler and three liquid nitrogen pumps are housed in the room with the helium cold box module in the Utility Building. The liquid nitrogen transfer lines are routed from the Utility Building to the detector along with the liquid helium transfer lines.

References:

1. "SDC Experimental Facilities User Requirements," SSCL Doc. # SCT-0000001.

13. Installation

13.1. Overview

A preliminary installation plan and schedule has been formulated for the installation of detector components in the underground hall. The installation schedule is given in Fig. 13-1. The entire detector installation period totals almost 43 months, with an additional six months for accelerator installation and commissioning. The installation process begins three months before beneficial occupancy of the interaction region hall. This chapter is a summary and update of previous and more detailed studies [1,2]. These studies have been used to size the underground hall and related shafts. See Chapter 12 for a description of the interaction region hall.

The critical path for installation is largely determined by the detector configuration. The barrel steel must be erected before the muon chambers can be attached and the barrel calorimeter inserted. The barrel calorimeter must be in place before the superconducting solenoid can be installed which in turn must be in place before the central tracking devices can be installed. The critical path is indicated in Fig. 13-1.

The installation process has been studied in Refs. 1 and 2 and for this report using three dimensional CAD modeling. This has allowed us to determine underground hall space requirements. Large fixtures and tooling necessary in the installation have been roughly engineered and modeled. Drawings of critical points in the installation sequence are given in this Chapter. A month-by-month simulation of the detector installation sequence is given in Ref. 2. Our philosophy has been to do as much detailed preassembly above ground as possible and to assemble below ground as large subsystems as possible. Section 13.3 contains a description of a detector repair scenario.

A detailed Assembly and Installation Plan is in preparation and will be updated regularly as the subsystem designs mature and as the detailed design of the underground hall progresses.

13.2. Installation sequence

13.2.1. Detector support system erection

This task begins prior to beneficial occupancy of the underground hall. In order to compress the time required for installation of the detector support system, components will be lowered from the surface to the hall floor in parallel with assembly and installation operations. The completed support system is approximately six meters high which necessitates temporary scaffolding to provide access to work areas during installation. The support system installation takes approximately four months to complete.

A system of base plates to support and distribute the detector load, carried by a system of hydraulic jacks, is carefully positioned, aligned, and grouted into place. Installation of the jacks can be initiated when the base plate installation has progressed sufficiently. When all 80 jacks have been set into place, their exact location is surveyed and the jacks are aligned. The hydraulic control and readout system is then connected and tested. The jacks are set to an approximately level alignment, ready to support the horizontal girders. The six horizontal girder sections are positioned atop the jacks. Three girder sections are bolted together on each double row of jacks to form one full girder on each side of the detector. Adjustments to the jacks may be necessary as the girders are installed to maintain a level alignment.

The 40 cross members are installed between the two horizontal girders and braced and cross-braced using threaded steel rods. The girders, with the attached cross members, are then aligned horizontally and vertically.

Six inclined plates are installed atop the horizontal girders. The plates are supported temporarily while a preassembled space truss is moved into place and bolted between the plates. Figure 13-2 shows the support system during the inclined plate installation.

After all inclined plates and trusses are installed, the load cell indicating system is installed and connected. This provides an indication of the sharing of load of the support structure among the hydraulic




FIG. 13-1. Detector installation schedule.

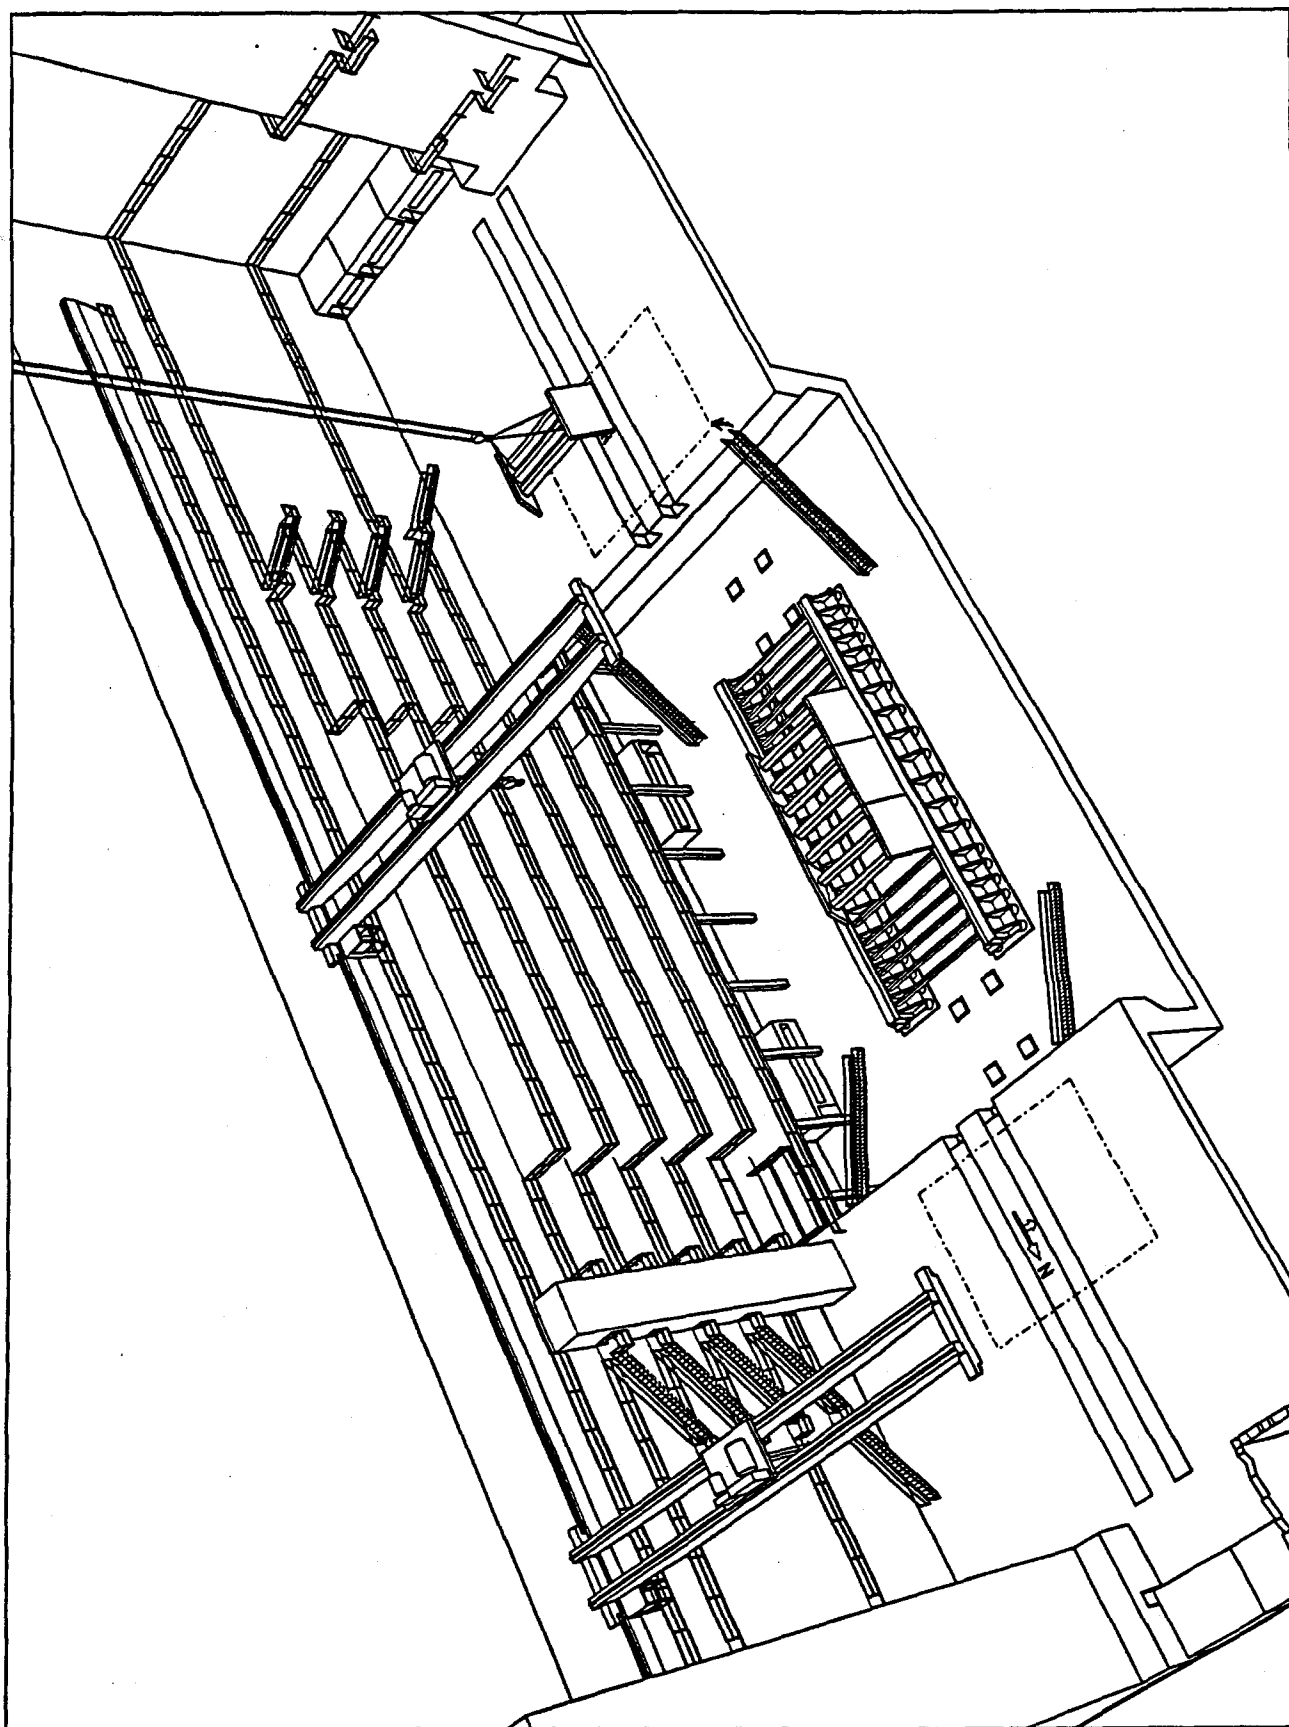


FIG. 13-2. Installation of detector support system in progress. Approximately month four of the installation schedule.

jacks. This, in turn, permits adjustments to the jacking system to distribute the load more evenly. A final torquing of all connecting bolts is performed after a final alignment of the support structure is completed. This completes the installation of the detector support structure.

Since the colliding proton beams in the interaction hall are at a slight angle with respect to the horizontal (approximately 2 mrad) the detector is tilted to align it with respect to the beams. Whether this tilt is accomplished during construction or after construction using the detector jacking system is under study.

13.2.2. Barrel steel erection

This task begins as soon as the support system is sufficiently rigid to begin supporting the weight of the barrel steel. The barrel steel erection occurs from February through August of 1996. Figure 13-3 presents an installation sequence for all 192 blocks. The block size and weight (approximately 85 metric tons) has been optimized with respect to installation procedure and crane capacities. Note that the short blocks of octants 2, 4, 6, and 8 are twice as thick as the long blocks of octants 1, 3, 5, and 7. Four short blocks and eight long blocks make up one of sixteen octagons that form the barrel toroid. See section 7.7 for design details of the barrel toroid. Octants are numbered in a clockwise fashion starting at the top looking south. The support structure on which the bottom octant rests is almost six meters above the hall floor level. The top octant rises 23 meters above the floor. A system of scaffolding and work platforms will be necessary to provide access during the installation process. The two bottom magnetizing coil halves for the bottom octant (octant 5) are placed into position on the top of the support system trusses using an overhead crane prior to beginning installation of the barrel steel. Lack of clearance prevents later insertion.

The steel blocks of the bottom octant are integrated with the support structure to increase overall structural soundness. The entire bottom octant is installed first, starting with a thicker block installed at approximately the longitudinal center of the detector. Additional blocks are then installed on alternating sides of the original block. When a sufficient number of blocks (probably three to five) have been installed such that the working distance between the northernmost block and southernmost will permit simultaneous operation of the two overhead bridge cranes, independent parallel installation at each end can commence. (This requires that either both assembly shafts be accessible for lowering the blocks or that some blocks must be pre-positioned in a temporary staging area on the hall floor before installation begins.)

Once the bottom octant is installed, some leveling and realignment of the support structure may be necessary. Next, the short thick blocks for octants 4 and 6 are installed. Again, the first blocks are installed near the longitudinal center of the detector. An octant 4 and octant 6 block are installed opposite each other for balance before continuing. Additional octant 4 and 6 pairs are added alternately north and south of the first pair until all rows have been installed. As with the bottom octant, simultaneous parallel installation on both the north and south ends can proceed once sufficient clearance is available to avoid interference between overhead cranes.

Installation of the long vertical blocks of octant 7 at the next stage of assembly creates a tipping moment which tends to destabilize the support structure. For this reason, as the blocks are installed, a system of temporary supports are bolted into place to counter this overturning moment. Installation of the vertical blocks for octant 7 starts at the middle and progresses toward both ends. Next, the long vertical blocks for the opposite side, octant 3, are installed. As above, the first block is installed in the middle, after which work progress outward toward both ends. A system of temporary supports is also necessary for octant 3 blocks.

Finally, installation of the top three octants completes the octagon. An assembly frame is used for the installation of the top three octants. The assembly frame supports and aligns the blocks of the top three octants during installation. First, the short thick blocks for octants 2 and 8 are installed starting at the middle. After the two short blocks are installed, the top, octant 1, long blocks (two are required to match the octant 2 and 8 short thick blocks) are then bolted into place atop the structure. Figure 13-4 shows the barrel steel at this stage of erection. A completed octagon is connected to its neighbors by side plates. The procedure is repeated, alternating north and south, until the toroid is complete.




FIG. 13-3. Barrel toroid steel installation sequence.




FIG. 13-4. Barrel steel assembly in progress showing the top three octants being installed. Approximately month ten of the installation schedule.

Although all blocks have now been installed, considerable effort remains to level and align the toroid structure to the required tolerances. Final torquing of all bolts used to connect the various blocks is not performed until the necessary alignment has been accomplished.

13.2.3. Barrel steel coil installation and Phase 1 barrel muon system installation

After the barrel steel has been erected, coils must be installed to magnetize the steel. This will take place from months 12 through 14. A test of the coils and a measurement of the magnetization in the barrel steel will follow in months 15 and 16. Installation of muon chamber support rails on the barrel and services such as gas and cabling proceeds in parallel with the installation and testing of the coils. Some muon chambers must be installed prior to the installation of the barrel calorimeter and associated cabling so that sight lines exist for their alignment. The installation of these chambers is referred to as phase 1 barrel muon system installation. Phase 1 muon chamber installation is scheduled to occur approximately between months 16 and 19. The installation sequence for this section is complex as it involves the interweaving of several activities vying for the same space. The main concern however is that the activities be completed before the barrel calorimeter is installed.

The installation of the coils involves the attachment to the barrel toroid of coil segments that are as long as the barrel. The handling of these segments requires careful rigging. Figure 13-5 shows a proposed method for the installation of a coil segment. A large amount of clear floor space in the hall is necessary for the handling of the coil segments. Some means for workers to reach the attachment points for the coils must be provided. This may be scaffolding or hydraulic lifts. The power connections and cooling water connections will also be made at this time. Other installation work may have to be curtailed during the barrel coil testing and field measurements.

The installation of the barrel muon chamber support rails and services extends from approximately month 12 through month 18. The use of scaffolding or hydraulic lift platforms necessary to complete this work is coordinated with the coil installation. Figure 13-5 shows a device on the operating floor ready to roll into the barrel steel and install a muon chamber support rail. Bridges that span the gap between operating floor and barrel steel are in place at this time. Once the rails have been installed they must be aligned. Phase 1 installation begins after muon barrel coil testing and field measurements are complete. All of the chambers for octant 5 (the bottom octant) and some of the chambers in octants 7 and 8 are installed in phase 1. At least one of the bridges that spans the detector pit is removed to allow octant 5 muon chamber installation. Both inner (BW1) and outer (BW2 and BW3) chambers of a tower will be installed at the same time for alignment purposes.

13.2.4. Forward toroid erection

The erection of the forward toroid steel takes place between approximately months 13 and 16. This task is not on the critical path. The forward toroids assembly is scheduled shortly after completion of the barrel toroid assembly to take advantage of steel erection crews and to complete heavy steel work as early as possible to provide a cleaner environment for later technical work. The north and south forward toroids and their muon systems and forward calorimeters must be ready for roll-in at month 43. The north and south forward toroids will be erected sequentially at their respective ends of the interaction hall. Each forward toroid system consist of two separate toroids (FT1 and FT2). Each toroid is made up of 16 wedge shaped blocks. Bracket plates with rollers support the toroids and allow movement of the toroids. Track plates are grouted into the hall floor to facilitate movement of the massive toroid systems.

FT1 and FT2 are assembled separately. The bracket plates with rollers are positioned first and then the bottom octant wedges are installed and bolted in place. The 14 other wedges are then bolted into place working on both sides of the toroid toward the top. During the assembly of the toroids prefabricated coils to magnetize the steel are looped over the appropriate wedges. Figure 13-6 shows the forward toroids being erected. After FT1 and FT2 are complete they are moved so that they can be bolted together using tie plates, spacers, and the support brackets. Support sleeves for the exterior forward muon chamber FW4 are then installed. An extension of this sleeve for FW5 is added after roll-in as is FW5. Supports for the




FIG. 13-5. Installation of barrel steel coils and muon chamber supports. Approximately month 13 of the installation schedule.

forward calorimeters and shielding for the forward chambers are installed while the forward toroids are on the operating floors.

13.2.5. Barrel calorimeter lowering and installation

The assembly of the halves of the barrel calorimeter will be completed on the surface and the halves will be separately lowered into the hall. Electronics and access platforms are attached to each barrel half before it is lowered into the hall. The barrel halves are lowered through the south assembly shaft onto the south operating floor level of the interaction hall. The weight of each half is approximately 1600 metric tons. A method for the lowering is presented in Refs. 1 and 2. Each barrel calorimeter half rests on a transport frame that is supported by wire ropes that in turn are supported by heavy duty jacks capable of lowering the weight. This is shown in Fig. 13-7. Once the calorimeter is on the floor of the interaction hall, the wire ropes are disconnected from the transport frames.

A central calorimeter support frame is installed within the barrel toroid prior to installing the barrel calorimeter halves. Each barrel calorimeter half is separately rolled into the barrel toroid using rollers mounted on the transport frames. The transport frames mate to the installed calorimeter support. Horizontal hydraulic jacks are used to accomplish the movement of the barrel calorimeter halves. The south bridge from the operating floor level of the hall to the barrel steel inner surface must be in place before a barrel half can be rolled in. Once both halves are within the barrel steel, they are aligned using jacking systems under each half and joined. The installation of cable ways associated with the calorimeter electronics then commences.

The lowering and installation of the calorimeter halves takes place from about month 18 through 21. The use of the floor space and the assembly shafts must be coordinated during this time.

13.2.6. Solenoid installation

The superconducting solenoid is lowered into the interaction hall after the barrel calorimeter halves are installed into the barrel steel. The solenoid is mounted on a mandrel which in turn is supported by a transport cradle. A monorail system is temporarily constructed in the muon barrel steel, extending onto the south operating floor, for the installation of the solenoid into the barrel calorimeter. An extension of the mandrel is added in the hall to aid in inserting the solenoid into the barrel calorimeter. The solenoid is moved along the monorail until the mandrel extension protrudes from the opposite side of the barrel calorimeter. The extension is then supported by a short monorail while the slide between solenoid and calorimeter is removed. The solenoid is then inserted. Figure 13-8 shows the solenoid ready to be inserted. Note that the monorail system occupies floor space under the south assembly shaft and limits use of the shaft. The cryogenics, power, and monitoring systems are hooked up after insertion. The insertion and connection occur approximately from month 23 to month 26. The solenoid is cooled and tested after insertion. The endcap calorimeters must be in operating position prior to full power testing of the solenoid as they form part of the flux return for the solenoid.

13.2.7. Endcap calorimeters lowering and solenoid field mapping

The endcap calorimeters are lowered at approximately month 25 and temporarily placed into operating position to permit full power testing of the superconducting solenoid. The endcap calorimeters are lowered in the same manner as the barrel calorimeter halves. The north endcap is lowered through the north assembly shaft and the south endcap through the south assembly shaft. These lowerings must be coordinated with the phase 2 barrel muon chamber installation (see below), which is proceeding at the same time. The endcaps are lowered with their electronics and access platforms attached. Figure 13-9 shows one endcap being lowered.

The field of the solenoid is fully mapped after testing of the solenoid. This occurs from month 30 to month 32. The endcap calorimeters must remain in position during mapping.




FIG. 13-6. Forward toroid steel erection. Approximately month 14 of the installation schedule.

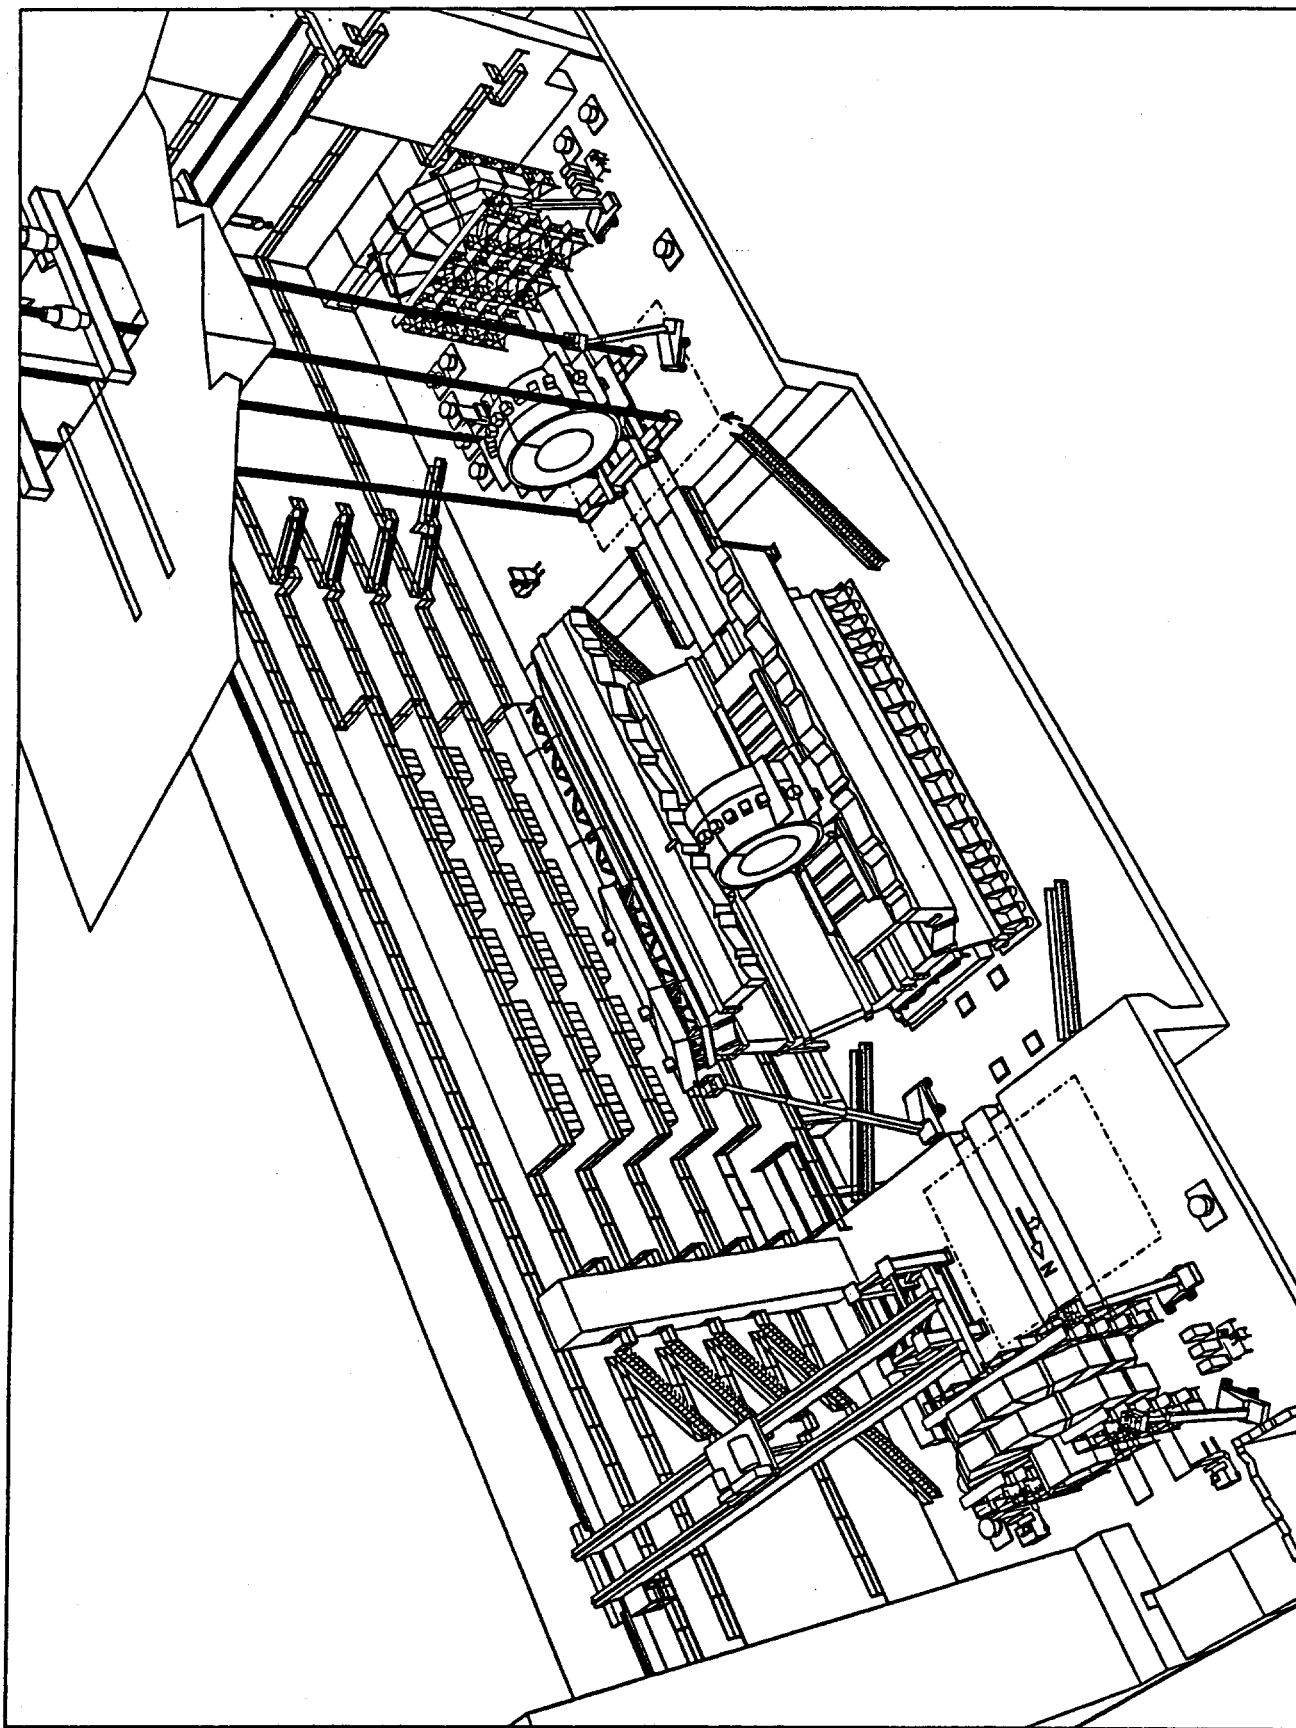


FIG. 13-7. View of interaction hall showing second barrel calorimeter half being lowered into the hall and first calorimeter half inserted in barrel steel. Approximately month 20 of the installation schedule.




FIG. 13-8. View of interaction hall showing the superconducting solenoid being installed. Approximately month 24 of the installation schedule.




FIG. 13-9. View of interaction hall showing an endcap calorimeter being lowered into the hall. Approximately month 25 of the installation schedule.

13.2.8. Phase 2 barrel muon system installation

Phase 2 barrel muon system installation is defined as the installation of all barrel muon chambers that do not need to be installed before the installation of the barrel calorimeter. Phase 2 is scheduled to begin after installation of the barrel calorimeter and the superconducting solenoid. The installation spans months 24 to 29. Phase 2 barrel muon installation is not on the critical path.

Figure 13-10 shows a BW2 and BW3 muon chamber assembly being installed. The concept depicted for installing the assemblies is to lower them completely assembled from the surface on a transfer frame that is mated to the end of the muon barrel steel allowing the package to be moved off the transfer frame and onto the support rails attached to the barrel steel. Simulations of this method have been done for all eight octants. The weight of the largest assembly is approximately 30 metric tons. Due to obstructions in some octants, such as the solenoid's cryogenic chimney, some assemblies can be installed from one end only of the barrel steel. Final alignment is done in position on the barrel steel. BW1 chambers are installed in sequence with projective BW2 and BW3 assemblies. This allows alignment to proceed in an orderly manner. The assemblies are lowered through both the north and south assembly shafts. Scheduling of the shafts and cranes is necessary to avoid conflicts.

13.2.9. Central tracking system installation

The central tracking system consist of the barrel tracking system, the silicon tracking system, and the north and south intermediate tracking systems. The installation and alignment of the central tracking system occurs approximately from month 33 to month 41. To install the central tracking system the endcap calorimeters must be moved from their operating positions onto the operating floors at each end of the hall. A monorail similar to that used to insert the solenoid is then erected. Work and equipment movement must be carefully scheduled during this time of high congestion. Figure 13-11 shows the hall during the installation of the outer tracking system.

The barrel tracking system is installed and aligned first. The silicon tracking system is then installed inside the barrel tracking system and aligned. The north and south intermediate tracking systems are installed following the silicon alignment. Cabling of all systems proceeds in parallel with their installation.

During the tracking system installation, installation of forward and intermediate muon chambers and the forward calorimeters are proceeding.

13.2.10. Endcap calorimeter installation

The endcap calorimeters are moved to their final operating positions after the central tracking system is installed and aligned. This occurs at approximately month 42. The final cabling and testing of the endcap calorimeters then begins. Segments of the beam pipe may be installed prior to installing the end caps so that connections within the calorimeter cavity may be made. Muon chamber shielding is connected to the rear of the endcap calorimeters before they are moved into the barrel steel.

13.2.11. Forward and intermediate muon systems installation

The installation of the forward and intermediate muon systems occurs in parallel with the critical path activities and after the completion of phase 2 muon barrel chamber installation. These activities occur from approximately month 30 to month 43 of the installation schedule. It is efficient to complete one type of chamber assembly on the surface before commencing assembly of a second chamber type. Therefore, the barrel chambers, are completed, then the forward muon chambers followed by the intermediate chambers.

The forward muon chamber system consists of chambers FW1, FW2, FW4, and FW5. Chambers FW1, FW2, and FW4 are installed on the north and south forward toroids when the toroids are on the operating floors. At least the bottom portions of chambers FW5 must be installed after the forward toroids are in operating position as the outer radius of FW5 extends below the bottom support rollers of the forward toroids. In this installation schedule we assume that all of the FW5 chambers are installed after the forward toroids are in operating position. Mounting fixtures and service lines are installed on the forward




FIG. 13-10. View of interaction hall during phase 2 muon chamber installation showing a BW2 and BW3 muon assembly being installed on the barrel steel. Approximately month 27 of the installation schedule.




FIG. 13-11. View of interaction hall showing installation of the barrel tracker. Approximately month 34 of the installation schedule.

toroids before chamber installation. This requires scaffolding or hydraulic lifts that will occupy operating floor space. Chambers FW2 are mounted on the inner faces of the forward toroids and chambers FW1 are mounted on frames extending inward from the FW2 chambers. These frames retract to move FW1 closer to FW2 to allow servicing of the endcaps and to allow the movement of the endcaps away from the interaction point to allow access to the central tracking system.

Figure 13-12 shows the installation of an FW1 segment. The method of installation depicted may not be the method finally chosen. A framework within the inner radius of FW4 is installed to aid in supporting FW4 and to support scintillation counter FS4 and chamber FW5 with FS5. The portion of this framework supporting FW5 and FS5 is installed after the forward toroids are in their operating positions.

The intermediate muon chambers are installed after the FW1, FW2, and FW4 forward chambers. The bottom octants of the north and south intermediate chambers can not be installed until the forward toroids have been moved into their operating positions and the bridges spanning the detector pit have been removed. The FW5 chambers installation must also wait until the removal of the pit bridges.

13.2.12. Forward calorimeter installation

The forward calorimeters are located within the FT2 forward toroids. The forward calorimeters are installed while the forward toroid systems are on the operating floors. There is flexibility in the scheduling of the forward calorimeter installation; it is assumed to occur after the installation of the FW4 chambers at approximately months 39 and 40. The forward calorimeters are lowered into the hall completely assembled. Each weighs approximately 150 metric tons. A jacking system similar to that used to lower the barrel and endcap calorimeters or a temporary crane with sufficient capacity is necessary to lower the forward calorimeters. Figure 13-13 shows the south forward calorimeter being lowered and the north forward calorimeter being installed. A forward calorimeter is positioned on a transfer frame which is aligned with support rails on the inner radius of FT2 and then moved into FT2 with hydraulic jacks. Alignment and service connections are then made.

13.2.13. Completed detector and accelerator components installation and commissioning

Figure 13-14 shows the detector in the final assembled state. The forward toroids have been positioned and the final forward and intermediate muon chambers installed. The detector installation is scheduled to be completed in spring of 1999. Accelerator component installation follows. Figure 13-14 shows the accelerator components installed. The installation of the accelerator components and the commissioning of the accelerator is assumed to take six months. The installation of accelerator components is presently under study.

13.3. Major repair or upgrade

Figure 13-15 shows the detector disassembled for major repair or upgrade. This may include a central tracking system upgrade or major endcap calorimeter scintillator replacement. In the figure, the forward toroids are retracted, the endcap calorimeters are moved onto the operating floors, and a monorail system has been erected for the removal of the barrel tracking system. Storage of accelerator components is shown within the hall. Some forward and intermediate muon chambers may be stored on the surface due to floor space limitations. Space under the assembly shafts will also have to be clear during lowering or raising operations. The figure shows that hall floor space will be entirely used and that careful planning is necessary for a tightly scheduled repair or upgrade period.




FIG. 13-12. Installation of an FW1 muon chamber on the north forward toroid system. Approximately month 35 of the installation schedule.

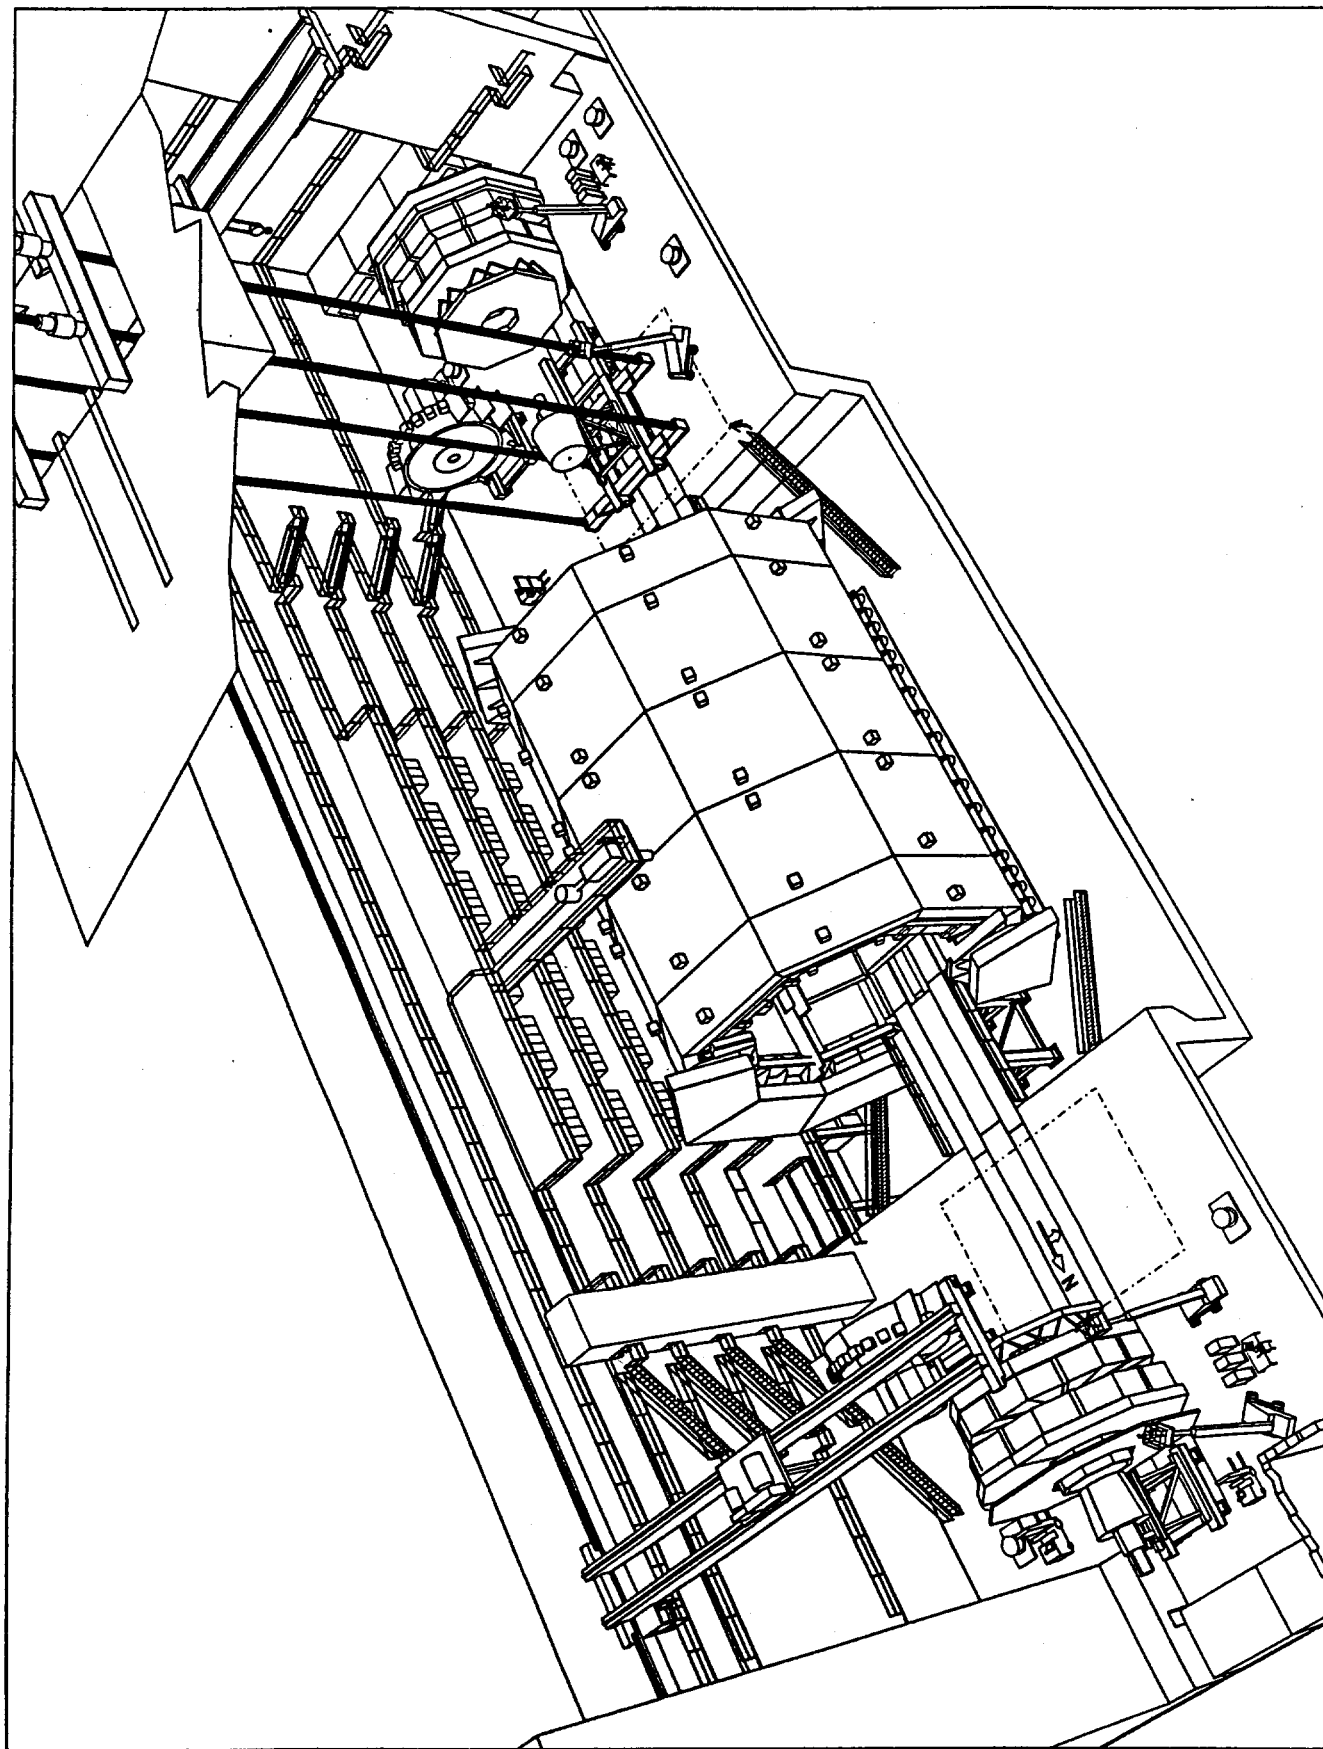


FIG. 13-13. View of interaction hall showing the lowering of the south forward calorimeter and the installation of the north forward calorimeter. Approximately month 39 of the installation schedule.

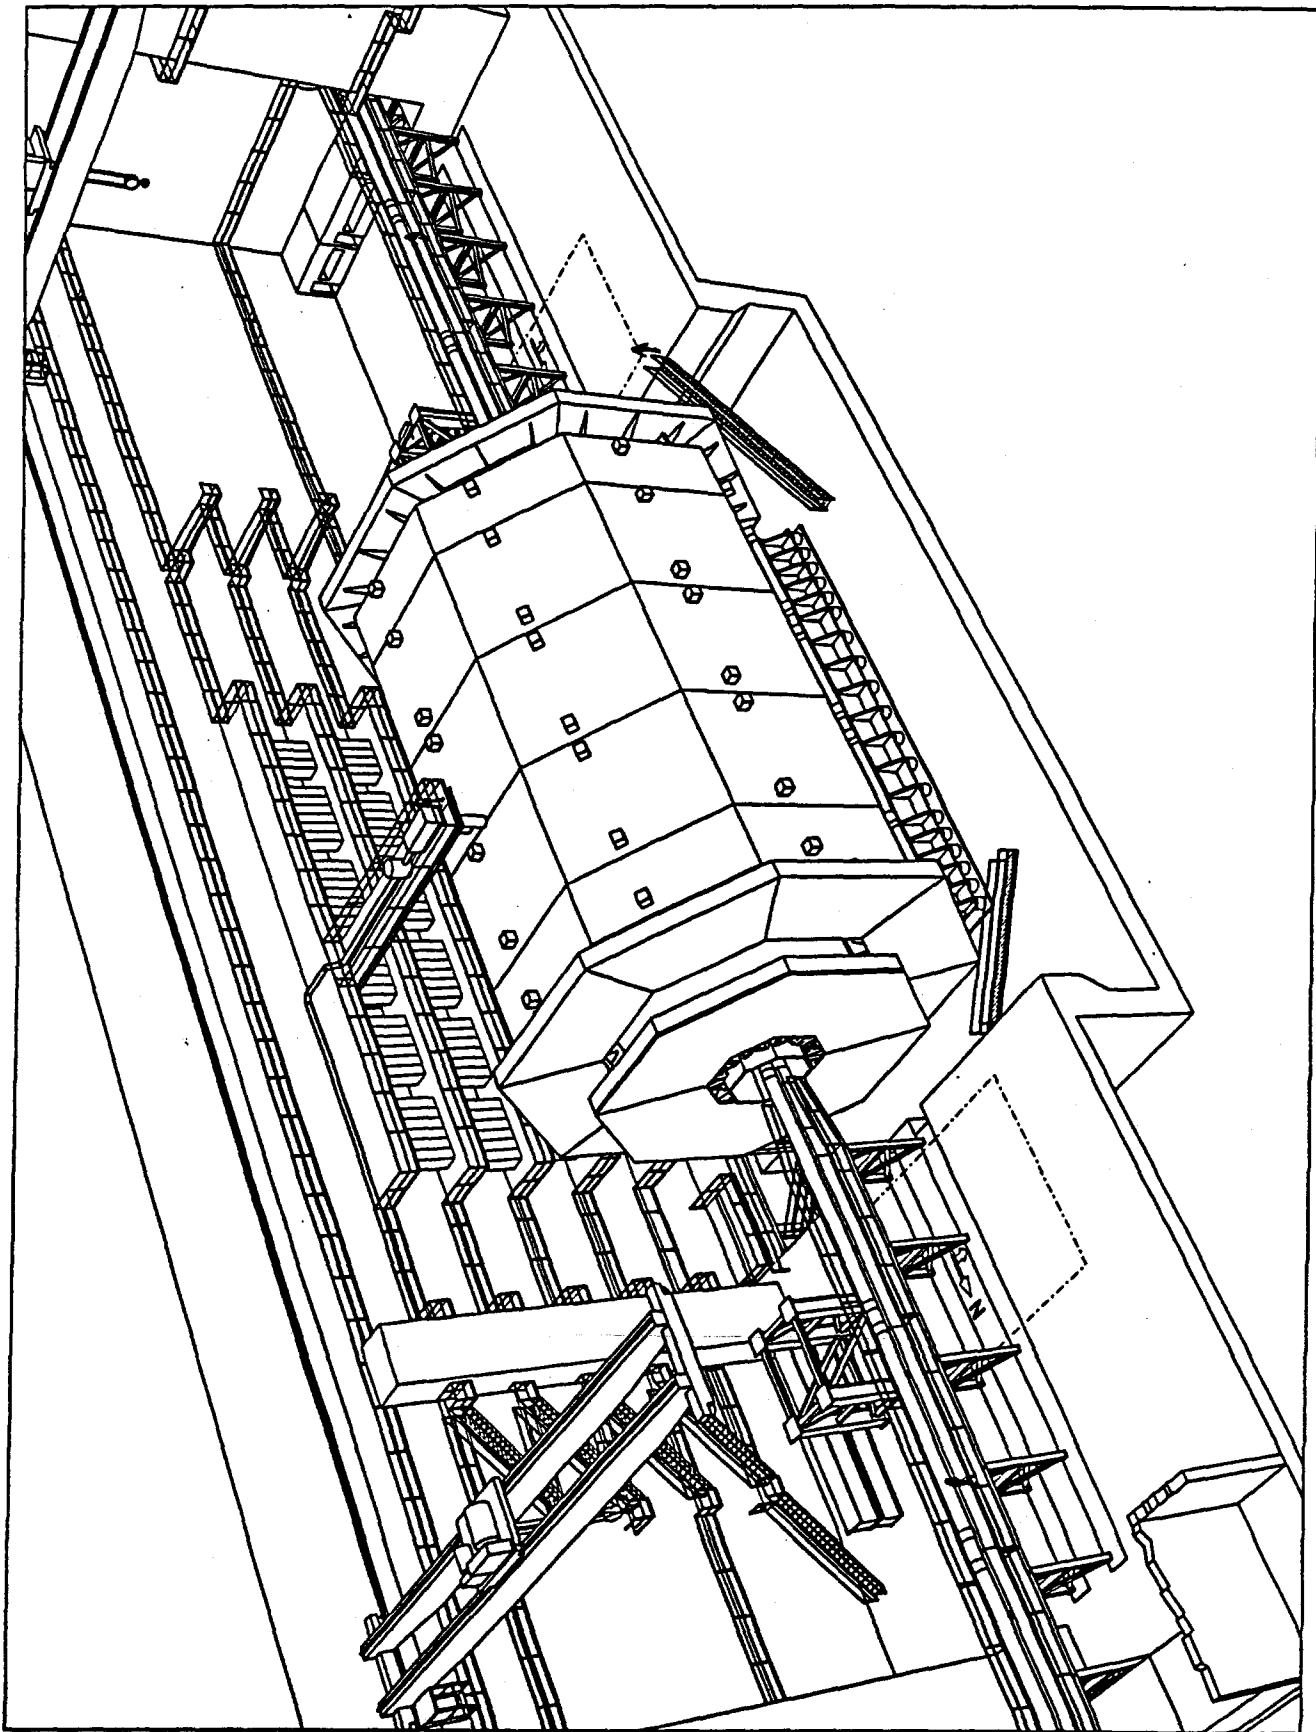


FIG. 13-14. View of the interaction hall showing the completed detector with accelerator components.
Month 45 of the installation schedule.

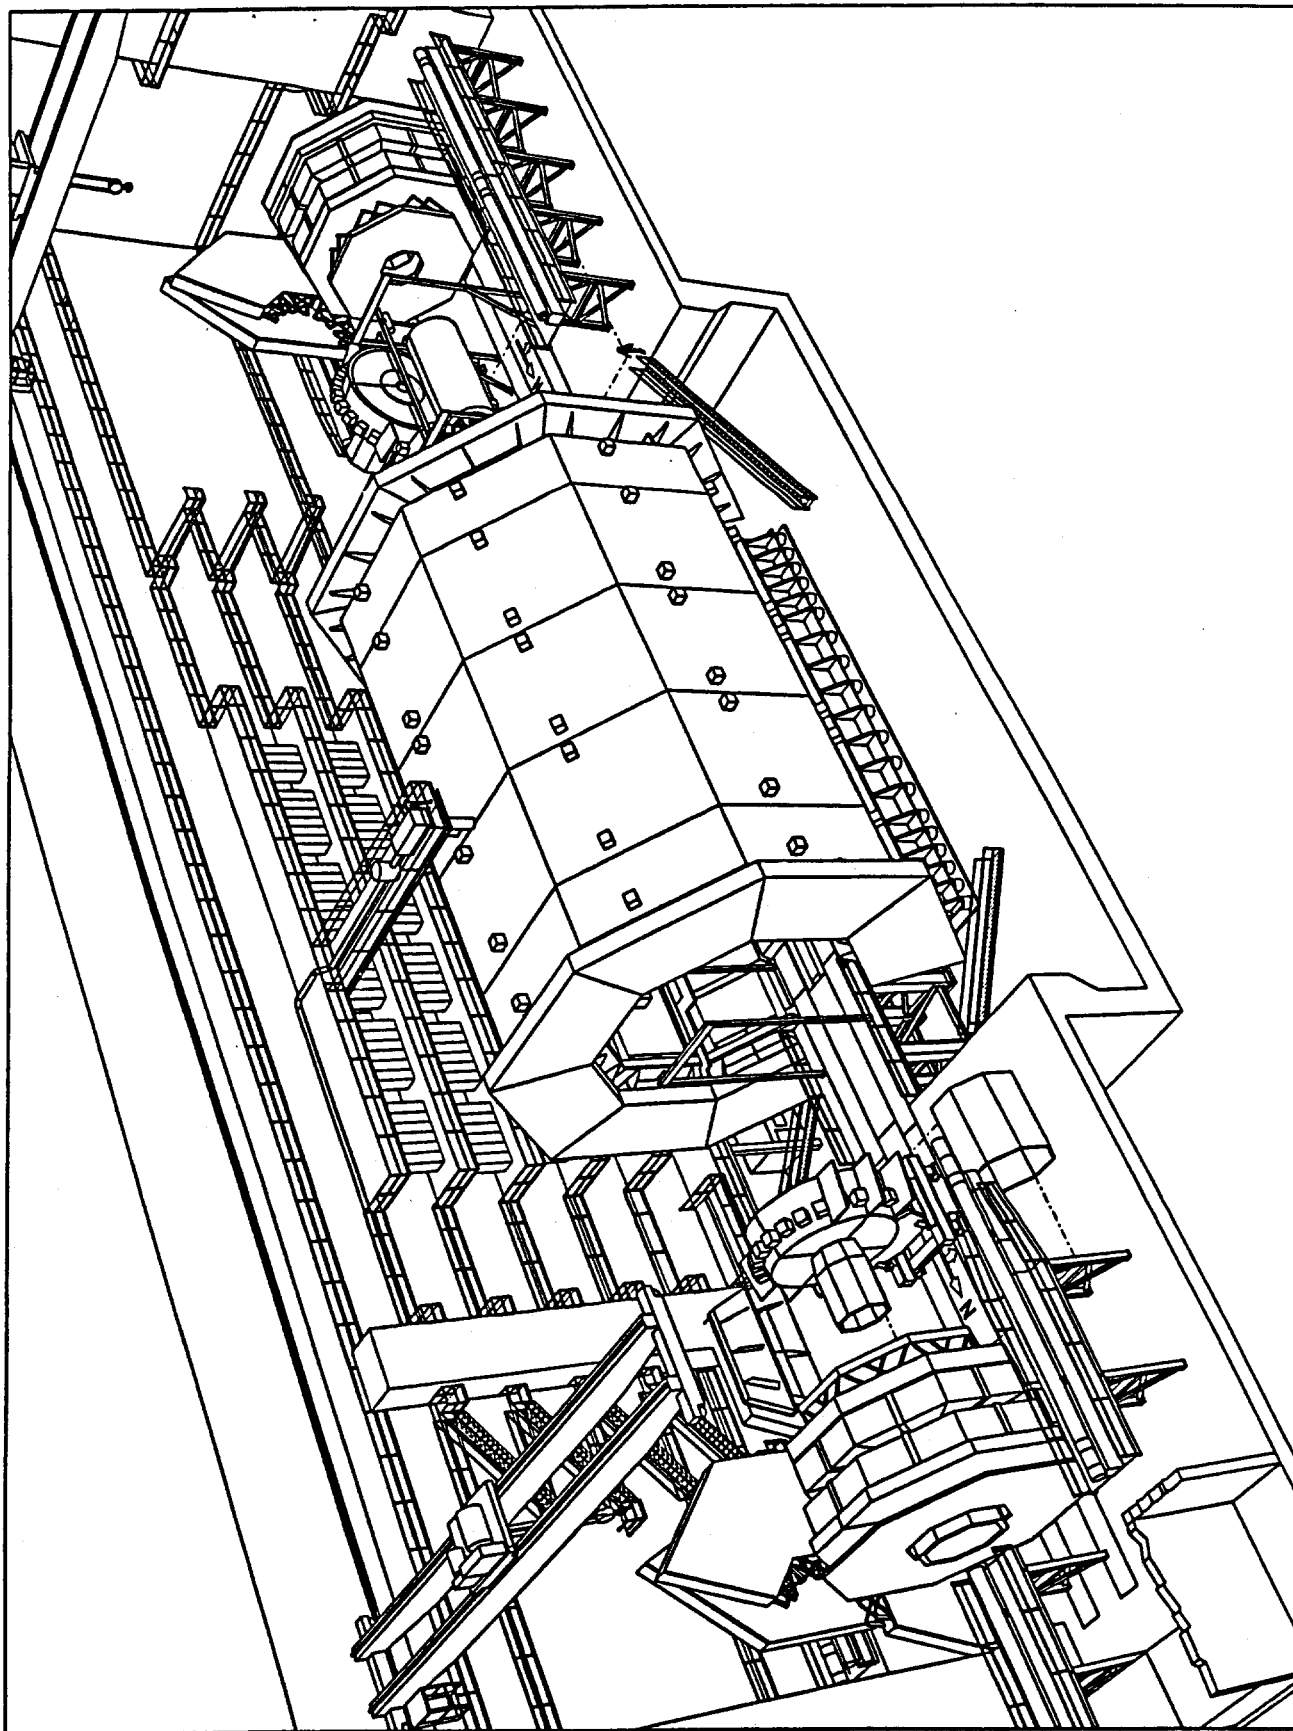


FIG. 13-15. View of the interaction hall showing the detector disassembled for upgrade or major repairs.

References:

1. A Design Study of Interaction Halls and Assembly Procedures for a Solenoidal Detector at the SSC, SSC Report No. SSCL-SR-1152 Rev. 1, April 1991.
2. Preliminary Construction Planning SDC Scintillator Calorimeter Detector, SDC Report No. SDC-91-117, October 1991.

14. Test and calibration beam plan

14.1. Introduction

The development of the detector subsystems requires extensive tests in particle beams of varying types and intensities. Once the design and prototype phase is completed, calibration beams will be required to evaluate and monitor the performance of detector subsystems, primarily calorimetry, over the entire lifetime of the experiment. The SDC intends to utilize both low and high energy beams at accelerator facilities throughout the world. Indeed, SDC beam tests have been completed in the last year or will be started in the next few months at accelerator facilities in the Peoples Republic of China, Japan, France, Russia and the United States (FNAL, BNL and LAMPF). We expect this broad utilization of international accelerator resources to continue during the design and fabrication phases of the detector project. Beginning in the last few years of this decade, we expect to fully utilize a dedicated SDC beamline at the test beam facility of the Medium Energy Booster at the SSC Laboratory.

In the sections below we briefly summarize our plans for test beam usage. The schedule for test beam usage is necessarily uncertain, since it depends critically on the availability of beam time at key accelerators for which definitive long range schedules do not now exist. Furthermore, the need for future test beam studies a few years hence depends on the results of tests in the near future, and thus the exact need cannot be precisely predicted. Nevertheless we present below our best estimate of the summary schedule for test and calibration needs during the construction of the detector.

14.2. Accelerator facilities

We expect to utilize the accelerator facilities listed in Table 14-1 for a large variety of test and calibration measurements. Low energy electron beams in the PRC, Japan and France have already been used to irradiate electromagnetic calorimeter modules. High intensity proton beams at LAMPF have been used to irradiate silicon detectors and electronics. High energy beams at Fermilab have been used to test electromagnetic calorimeter modules and to test hadronic response for different absorber configurations. Beams at Brookhaven have been used for SDC calorimetry tests. Fermilab beams have been used to test scintillating fiber tracking and for muon studies. The muon effort, in particular, benefitted from the use of a high energy muon beam, since the issue of electromagnetic showers accompanying energetic muons is of primary importance for muon tracking.

In the future we expect these activities to expand, and dedicated facilities will be required, first at Fermilab and finally at the SSCL Laboratory.

Major facilities at Fermilab and eventually at the SSCL will be the primary focal points for test beam and later calibration beam work. Calorimeter test and calibration requires high energy beams that can only be provided now by Fermilab and CERN and later by the SSCL. We have chosen to base our operations in the next few years at Fermilab. Since much of the calorimetry construction will take place at Fermilab and ANL, Fermilab is the natural test location. We expect that much of our other test beam needs can also be satisfied by sharing the beamline used for calorimetry tests and calibration at Fermilab.

Our present construction schedule requires the central calorimeter to be completed and ready for installation by late 1997, which means that calorimeter modules must be tested well before this time. The present SSCL schedule shows that the test beam facility from the 200 GeV Medium Energy Booster will be available in early 1997. This is too late to be useful in the testing of production central calorimeter modules. If the MEB schedule can be advanced by one year, then some testing of modules could occur at the SSCL. In any case it is our intention to fully utilize at least one dedicated beamline in this facility for calibration and monitoring of the central calorimeter response. In addition, the SSCL beamline would be used to check the response of tracking and muon system components as required. Finally, upgrades to the detector can be expected and we will use the SSCL beams to test prototypes for this purpose.

Table 14-1
Accelerator facilities that the SDC has utilized or plans to utilize
during detector design and construction.

Facility	Country	Major area(s) of utilization
Fermilab	United States	Calorimeter tests, muon system response and tracking system tests
Brookhaven AGS	United States	Calorimeter tests, muon chamber tests and tracking system tests
LAMPF	United States	Radiation damage and silicon detector tests
KEK	Japan	Radiation damage tests (electron beams), tests of silicon detectors and tracking systems
IHEP Beijing	China	Radiation damage tests(electron beams)
LURE Orsay	France	Radiation damage tests(electron beams)
CERN	Switzerland	Tracking system tests, calorimeter module tests
IHEP Protvino	Russia	Muon chamber tests

14.3. Test and calibration beam summary schedule

In Fig. 14-1 we present a summary schedule of test and calibration needs by subsystem for the remainder of the decade. The schedule corresponds to the SDC detector subsystem development and calibration needs and may not correspond to the availability of beam time at the accelerator facilities. The test beam schedule is closely related to the schedule for construction of the detector and its installation. A summary schedule for detector construction is given in Chapter 15 and the installation schedule and plan is described in Chapter 13.

14.4. Detector subsystem tests and calibration

In the sections below we briefly describe the anticipated test or calibration needs for each detector subsystem.

14.4.1. Silicon tracker

The development of the silicon tracker will be closely coupled with radiation hardness tests in a continuation of our ongoing three year old program at LAMPF at the Los Alamos National Laboratory. Beam tests to confirm and monitor the resolution of modules consisting of double-sided detectors, front-end electronics and readout systems will also occur.

The proton irradiations will consist of one or two week long runs per year during the running cycle of LAMPF, typically during the summer, from FY1992 to FY1995. With these irradiations, we will be able to check the radiation hardness of the double-sided silicon detectors, and the bipolar and CMOS front-end electronics at various stages of development and production.

In the prototype phase (FY1992-FY1993), beam tests in a KEK pion beam will be used to measure the resolution of the silicon detectors with the full front-end electronics chain, both before and after irradiation. In the preproduction phase (FY1994-FY1995), modules will be tested with emphasis on system issues. Two runs per year at KEK will be sufficient.

For the production verification/calibration phase (FY1996-1999) we do not foresee the need for beam tests, but will rely on laser/X-ray calibration systems to do the calibration and alignment tasks.

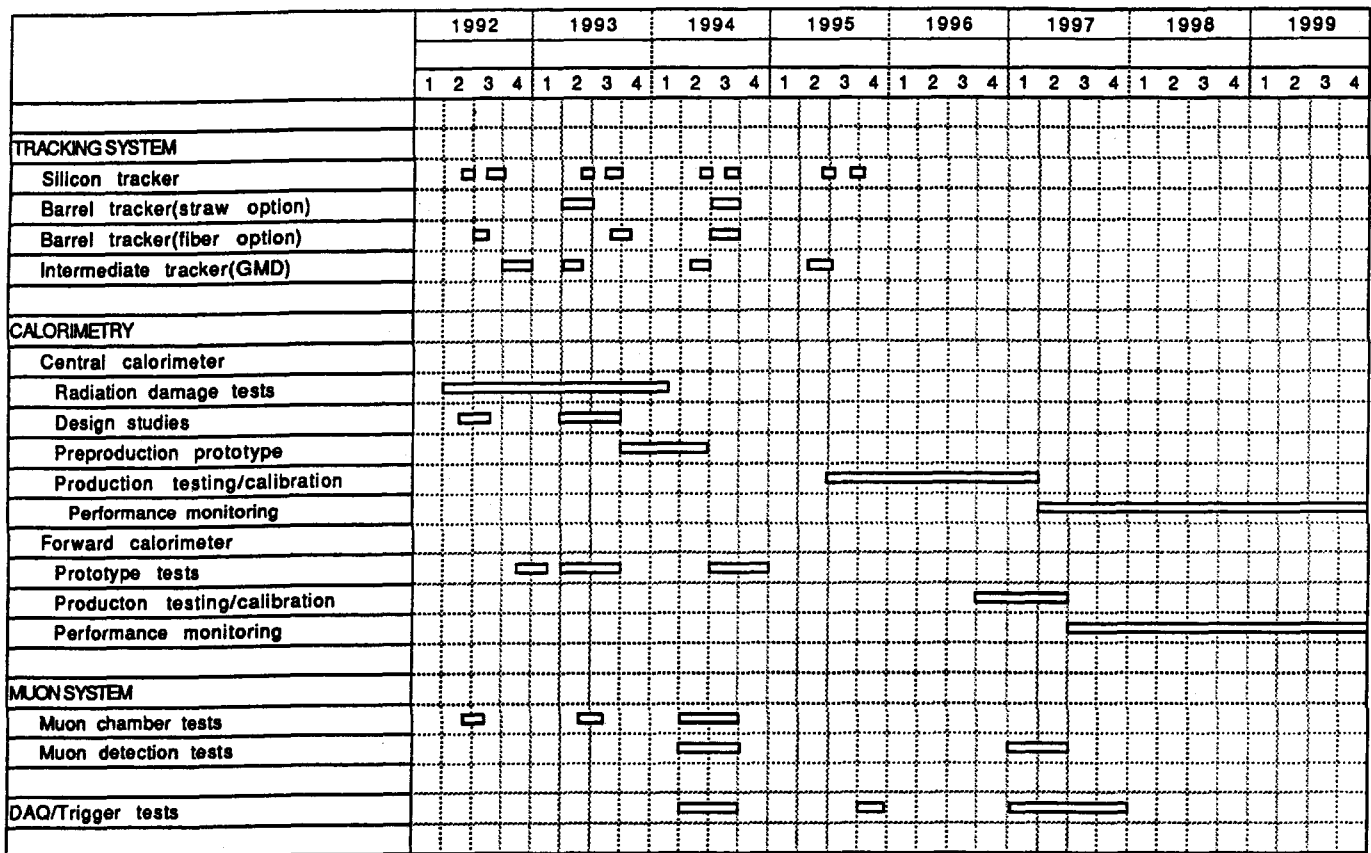


FIG. 14-1. Summary of the preliminary schedule for test and calibration beam needs for detector subsystems.

14.4.2. Straw-tube tracker

Most of the testing of the straw-tube tracker prototypes and modules can be done with cosmic rays. However, we plan to assemble a small number of prototype modules in a three or four layer arrangement for a systems test in 1993. This test would be done with prototype electronics and would be used to verify the system response, measure resolution and possibly high rate capability, depending upon the availability of a suitable beam. This test could be done at Brookhaven or possibly at Fermilab or KEK. If necessary, additional high rate tests would be performed in 1994–1995 at Fermilab to verify the performance of the modules, including trigger performance.

14.4.3. Gas microstrip intermediate tracker

Beam tests of prototype gas microstrip tiles are planned at CERN in the second half of 1992 and in 1993. These tests are aimed at determining the pulse shape, time and spatial resolution, angle dependence, *etc.* of the prototype tiles fabricated in the United Kingdom or Canada. The test beam plan beyond 1993 depends upon the results of the initial tests. However, a systems level test may occur in 1994 and 1995 as shown in Fig. 14-1.

14.4.4. Scintillating fiber option

A sequence of beam tests is planned for the scintillating fiber tracker.

Radiation Damage Tests.

While many radiation damage tests can be carried out using radioactive sources, there is always some uncertainty concerning the validity of such tests. Therefore, we plan to carry out radiation damage tests at Fermilab in the collider tunnel similar to the test that has just been completed, T-851. The environment within the tunnel is similar to that expected at the SSC. We expect these tests to be conducted almost continuously over the next 5 years, until 1997.

Other tests will include detailed radiation damage studies of the various elements of the tracker, such as the composite material of the support structure, glues for holding the fibers in place and the fibers themselves. These will be small studies and many can be carried out using sources. However, we expect that a number of small scale beam tests will be required. We may use a test beam at Cornell at the positron target starting in 1993 and again in 1995. There is a possibility of work in 1994 at KEK in Japan.

System Verification Tests

Table 14-2 lists the major tests that will be performed.

Table 14-2
Summary of the tests planned for the scintillating fiber option.

Test	Channels	Location	Year(s)	Beam
1	128	BNL	1992	10 GeV/c π
2	1000	FNAL/SSCL	1993	Cosmic rays
3	256	BNL	93/94	10 GeV/c π with magnet BM109
4	10 K	FNAL	94/95	To be determined

The tests are:

1. A beam test will be carried out at BNL starting in May 1992 using a pion beam with energies up to 10 GeV/c. The purpose is a proof-of-principle test, using 128 channels of Hist III VLPC's. The photoelectron yield and resolution will be measured for 3 superlayers.
2. A thousand channel test will be carried out using cosmic rays at FNAL early in 1993. The purpose of this test is to show stable long term operation of a VLPC-scintillating-fiber system. Operational problems that show up will be corrected as part of the final system design.
3. Improvements found while operating test 2 will be incorporated into test 3. This smaller scale test will be done using a magnet at BNL. A magnet the size of a BM109 or larger should be adequate for this test. The trigger will be studied.
4. A 10,000 channel test is required in order to test a full cryostat of VLPC's and the DAQ system. The full cryostat is the unit cell of the fiber tracking system. The full fiber tracking system is simply the replication of this element. The FNAL test beam should be available near the end of 1994 and the beginning of 1995.

14.4.5. Central calorimetry

Beam tests for the central calorimeter may be divided into four categories: (1) exposures of scintillators and small electromagnetic test modules to intense electron beams to measure radiation damage; (2) preproduction tests of prototypes to assess design; (3) testing and calibration of production modules; and (4) continued calibration of modules to monitor long-term performance after installation of the central calorimeter. Each of these areas is discussed briefly below.

Central calorimetry radiation damage beam tests

Previous efforts by SDC groups, at Orsay, KEK, and Beijing have concentrated on providing an existence proof for the survivability of the barrel calorimeter. Exposures of about 1 Mrad of low energy electrons were used to establish this. Doses appropriate to the endcap region will be studied in the future. Tests will continue at KEK and Beijing up to the level of 100 Mrad. This dose represents more than 100-year operation at $|\eta| = 3$ at the design luminosity. In 1993, tests will shift to the study of new scintillator materials developed at Bicron, Kuraray and Fermilab. It is anticipated that electron beam tests will continue through at least 1993.

Central calorimetry prototype tests

During 1991, beams at Fermilab were used to test electromagnetic modules, shower maximum detector concepts, and hadronic response. In 1992, plans are underway to test EM design concepts for the barrel and endcap at BNL. The current plan is to use the results of this series of tests to define a preproduction prototype in the middle of 1992, assemble a full scale barrel wedge prototype (~40 tons), and test it in a beam at Fermilab in the latter half of 1993. This aggressive schedule of tests is needed if the production schedule for calorimeter modules is to be met.

Central calorimetry production testing and calibration

Much of the final assembly of the calorimeter modules will be based at FNAL and ANL. It is therefore natural that tests of these production modules be made at Fermilab. The precise number of modules requiring test remains to be determined, but some subset of the modules must be beam tested as a final quality check. High energy beams are required. For example, a 200 GeV electron beam would test the module uniformity and energy resolution at about the 1% level.

Barrel modules from the production assembly will begin to appear in 1995 and will be tested. Endcap modules will undergo similar testing beginning in 1996.

Central calorimetry-performance monitoring

Long term response of selected modules will be monitored in the test beam derived from the Medium Energy Booster at the SSC Laboratory starting in 1997. Various configurations of modules (cracks, edges, etc.) will be studied and calibrated.

14.4.6. Forward calorimetry

The test beam program for the forward calorimeter will proceed in three phases; R&D, prototypes, and calibration.

Before a technology choice is made, both the high pressure gas and the liquid scintillator concepts will be studied using a series of small modules. Studies have already begun using small electromagnetic modules, and will continue with larger hadronic prototypes.

Different types of modules must be tested for response. Energy resolution, linearity, the effect of incidence angle and other features must be measured in high energy electron and hadron beams. The high pressure gas group intends to test a hadronic module at CERN or Brookhaven in late 1992 or early 1993. The liquid scintillator group plans a series of tests at CERN beginning in the latter half of 1992, starting with small electromagnetic modules, and progressing to hadronic modules in 1993.

Radiation damage tests must also be done, and a subset of the modules will be heavily irradiated. If necessary, these tests can be done using intense gamma ray and neutron sources, without using valuable accelerator test beam time.

Subsequent to the final choice of technology, several full-scale prototype modules will be constructed, and subjected to tests in a high energy hadron beam in 1994 or 1995. As modules of the final calorimeter are constructed, it will be necessary to calibrate at least a subset in a hadron beam. This will occur from 1996 onward. Even after installation is completed, beam studies will continue on the full scale prototypes to check the long term stability of the calorimeter response, and to understand long term irradiation effects. These will occur at the SSCL.

14.4.7. Muon system

Beam tests will be required to measure and evaluate the performance of the drift tubes and associated electronics, the trigger scintillation counters, and possibly the trigger Cerenkov counter modules.

Drift tube tests will take place in 1992 at BNL and IHEP (Protvino). The 1992 tests will measure only short tubes about 1 m long, and will involve the round field-shaped drift tubes that are the basis of the barrel, intermediate and forward muon system conceptual designs. The design specifications require that these tubes have 250 micron single-track precision and 5 mm two-track resolution. Several sizes of tubes will be tested, ranging in diameter from 4.2 cm to 9.0 cm; Several field-shaped, rectangular cells will also be tested and compared in performance to the baseline design. Continued calibration of modules to monitor long-term performance and ensure uniformity of production will be necessary. Long term radiation damage assessments will also have to be performed, but likely can be done with radioactive sources.

In 1993-1994, tests of full size drift tubes (9 m in length) will occur. The test facility must have the capability of full transverse motion for these tubes. At BNL, the test beam enclosure would require modification to accommodate such transverse motion. The possibility of conducting such tests elsewhere will be explored.

During 1995, tests can continue at Fermilab. In the period 1997-1999 tests of selected small modules consisting of many glued tubes (such as IW2) will take place at the SSCL.

The scintillation counters that will be part of the muon trigger will also require evaluation in a test beam. The Cerenkov counter design and performance may have to be verified in a test beam if it is decided to proceed along this upgrade path.